

RL-TR-97-46
In-House Report
August 1997



TOWARD AUTOMATING INTELLIGENT INFORMATION PROCESSING

Paul S. Yaworsky

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

DTIC QUALITY INSPECTED 2

Rome Laboratory
Air Force Materiel Command
Rome, New York

19970821 132

Although this report references limited documents(*), listed on page 60, no limited information has been extracted.

This report has been reviewed by the Rome Laboratory Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

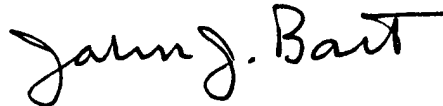
RL-TR-97-46 has been reviewed and is approved for publication.

APPROVED:



CARMEN J. LUVERA
Chief, Electronics Systems Engineering Division
Electromagnetics & Reliability Directorate

FOR THE COMMANDER:



JOHN J. BART, Chief Scientist
Reliability Sciences

If your address has changed or if you wish to be removed from the Rome Laboratory mailing list, or if the addressee is no longer employed by your organization, please notify Rome Laboratory/ERST, Rome, NY 13441. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</small>				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE August 1997		3. REPORT TYPE AND DATES COVERED In-House January 92 - June 95
4. TITLE AND SUBTITLE TOWARD AUTOMATING INTELLIGENT INFORMATION PROCESSING			5. FUNDING NUMBERS PE - 62702F PR - 2338 TA - 02 WU- 5D	
6. AUTHOR(S) Paul S. Yaworsky				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Rome Laboratory/ERST 525 Brooks Rd. Rome, NY 13441-4505			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Rome Laboratory/ERST 525 Brooks Rd. Rome, NY 13441-4505			10. SPONSORING/MONITORING AGENCY REPORT NUMBER RL-TR-97-46	
11. SUPPLEMENTARY NOTES Rome Laboratory Project Engineer: Paul S. Yaworsky/ERST, (315) 330-3690				
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) In order to embed intelligence into automated systems, we must be able to identify basic principles upon which an intelligent information processing model can be constructed. This means that we must be able to describe the "process" of intelligence. However, the foundations of intelligence do not exist in a tangible form. As a step towards addressing this unwieldy situation, we discuss various related topics involving several disciplines. These and other important concepts must be considered in concert if intelligence is ever to be automated. Artificial intelligence and neural network technologies hold much promise for the automation of intelligence. These technologies are also relevant to various missions at Rome Laboratory. We discuss how the automation of intelligent information processing impacts the methods of reliability science and electromagnetic technologies.				
14. SUBJECT TERMS intelligence, automation, process, order, signals, change, information processing, artificial intelligence, neural networks, reliability, electromagnetics			15. NUMBER OF PAGES 72	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

CONTENTS

0. EXECUTIVE SUMMARY	1
1. INTRODUCTION	3
1.1 Order, Organization and Information	4
1.2 Change	5
1.3 The Nature of Data	5
2. THE DATA SPECTRUM AND THE OCTAVE RULE	7
2.1 Describing an Elusive Process	7
2.2 The Data Spectrum	7
2.3 The Octave Rule	9
3. ARTIFICIAL INTELLIGENCE AND NEURAL NETWORKS	11
3.1 Attempts to Automate Intelligence	11
3.2 Lessons Learned	12
3.3 Future Directions	15
4. DYNAMICS AND COUPLED OSCILLATORS	16
4.1 The Nature of Changing Data	16
4.2 Oscillators	17
4.3 State Attractors	18
4.4 Modeling Dynamic Behavior	19
5. INTELLIGENT INFORMATION PROCESSING	22
5.1 The Quest for Knowledge	22
5.2 Consciousness	23
5.3 Goals	25
5.4 Decisions	26
5.5 Learning	27
5.6 Generalization	29

5.7 Common Sense	31
5.8 Intelligence and Understanding	32
5.9 The Importance of Time, Scale and Form	33
6. ON AUTOMATING INTELLIGENT INFORMATION PROCESSING	36
6.1 Describe the “Process”	37
6.2 Identify Sub-Processes	38
6.3 Model Key Components (the “Main Points”)	39
6.4 What Really Matters	41
6.5 Guidelines	42
6.6 A Word of Caution	46
7. IMPACT ON RELIABILITY SCIENCE AND ELECTROMAGNETICS TECHNOLOGY	47
7.1 Signal Activity	48
7.2 Physical and Mathematical Concerns	50
7.3 Manual Tools of the Trade	52
7.4 Automating Tools of the Trade	53
7.5 Adding Intelligence to Automated Tools	54
8. SUMMARY	56
9. FINAL REMARKS	57
10. REFERENCES	58

Relevance to the Air Force

We are concerned with establishing a framework for an intelligent information processing paradigm to be used with automated systems within the US Air Force (USAF). This framework involves a process, one which is complex, widespread and extremely important. The process involves intelligence -- the functions produced and enabled by the brain. We envision this process as a series of transformations which takes many low level, time-sensitive signals and produces fewer but higher level, more time-insensitive constructs. Our goal is to automate portions of this process and be able to integrate this functionality into Air Force systems.

The USAF constantly investigates ways to help position itself in the "Information Age." For example, in 1995 the Air Force Scientific Advisory Board initiated a study called New World Vistas, the goal of which was to assess and predict how future science and technology would influence military operations. The results of this study show "information" to be a core concept. Air Force 2025, a thirty year look into the future done by Air University, identified concepts, capabilities and technologies needed for the US Air Force to remain the best in the world. Again, the concept of "information" was a significant component in the study. Finally, the Air Force Core Competencies, which define the strategic vision to take us into the next century, were recently modified to include "Information Superiority." Obviously we in the Air Force cannot address all the issues regarding information, but we cannot afford to pass up the opportunities which related emerging technologies present to us.

The need for this work is evidenced by many factors. Technology is changing at an accelerated rate. Management of information in one form or another is a significant aspect of the USAF mission. We have to capture data, harness information, create knowledge and make decisions, all on a daily basis. If we do not do this well, for instance by using inappropriate information or by making decisions at an improper level, we will have to pay the consequences. Not only must we be able to process information intelligently but we must also do so efficiently. This report addresses how the computer may be used to help accomplish this task. Never before has the technology been available. Now that it is becoming so, we must develop appropriate models. This report and the work that flows from it are aimed at developing more efficient ways in which the Air Force can excel in this information-driven world.

Acknowledgments

The work described in this report was performed through close collaboration with James M. Vaccaro. Many of the ideas and concepts presented here would not have come to fruition without his insight, openness and technical expertise.

I thank all who reviewed and commented on this material, with a special thanks to Kenneth Siarkiewicz for his assistance and encouragement.

This work was supported in part by the Air Force Office of Scientific Research (AFOSR 93-2304H101).

0. Executive Summary

In this report we are concerned with trying to establish a framework for modeling a particular kind of process. This process involves intelligence and the intelligent processing of information. Additionally, we are interested in discovering ways to automate at least portions of this process. Thus our overall goal is to automate intelligent information processing. Finally, as part of our mission, we are concerned with how this automation may impact the related disciplines of reliability science and electromagnetics. Given the magnitude of this task and the scope of this work, this report addresses these topics at a high level.

Any discussion of a topic such as intelligent information processing will no doubt encounter problems with, among other things, terminology and the meaning of words. We realize that the terminology used and meanings given or implied here may be different than yours. We hope that you will at least recognize the main points we try to get across. We must also mention that there is no one way to describe this process. In fact, the process is so complex that it is for the most part a mystery. Thus our description here is really just a beginning.

So why talk about this process at all, if it truly is a mystery? Why is it so important? What does this process do, and where is it going? First of all, we all process information, all day long, every day, with varying levels of intelligence. Most of this is for our own personal benefit. But when we are involved with developing complex systems which involve the use of computers in some manner or fashion, we are forced to draw a line somewhere between what we do personally and what we may have our computers do for us. The line between the two is unclear and the distinction is becoming more and more important.

Some major goals of intelligent information processing are to reduce uncertainty, to gain knowledge, and to establish truth. All these feed our need and ability to make intelligent decisions. All of this constitutes a process, which admittedly involves many sub-processes. If we are to automate any portion of these processes and apply them to specific domains, we must first realize the fundamental principles involved. Without knowledge of the underlying constructs, we cannot expect to build anything solid. Unfortunately, a cohesive, concerted framework does not currently exist for intelligent information processing. Artificial intelligence and other related fields have made many inroads, but alas -- they have not cracked the nut!

What is needed is a model which will support the functionality needed to accomplish the task at hand. Our model initially consists of ideas, described in this report, which will eventually be tested by scientific method. To realize the benefits of this model, obviously much more work must be done to develop appropriate details that go along with the best top level constructs available. Once we have the shell of a model, we can then further develop and refine the model over time. When the model is ready to be applied, we will do so and see how well it performs on actual problems.

Many ideas are spread throughout this report. We attempt to tie them together in some kind of cohesive fashion. We emphasize the word attempt, because when it comes to information processing and intelligence, it is not easy to make sense of it all. Our observations are based on research, compilation and correlation of our work and the related work of many independent researchers. We bring up simple concepts, yet they are undeveloped with respect to automation. These concepts are pieces of an enormously

complex puzzle. This is not a rigorous, detailed analysis. Here we offer a different perspective, a general evaluation of the state of the art, so to speak, a common sense approach which should allow us to better automate portions of this complicated process. We make what we believe are some interesting observations, and intend to open this topic up for further discussion.

1. Introduction

A basic premise of this work is that structure and order exist within our world, and this fact should directly impact our attempts to automate intelligent information processing. We can and should exploit the natural order present. Our brain holds a model of the external world, or the physical world around us, but the brain's internal model is more than just a physical representation. The mental processes which represent and approximate reality in its many forms have not yet been revealed. We attempt to identify fundamental processes. We emphasize that with respect to intelligence and automating intelligent information processing, we will not try to exactly match the abilities of the human brain, but rather to match only certain processes or functions. The specific functions we choose to model or emulate are yet to be determined, but no doubt we will somehow have to exploit the natural order present.

The ultimate goal is to process data intelligently. This occurs naturally in the brain. The goal of our work is to get computers to do some of this. Our brain is a user of these so-called intelligent processes. It is also a developer, with self-organization built in. This occurs naturally. What is not so natural is to develop intelligence in a machine (i.e., a computer). As evidenced by the history of artificial intelligence, it has taken a tremendous amount of effort to come as far as we have today, and we still lack a solid framework on which to build intelligent machines.

When we mention the words "our model" in this report we actually mean an abstract model. Since we must initially talk about a framework of a model, it is difficult to talk about a "real model" before it is built. We will stress the importance of order and organization to our model, and how information processing fits into this model. While order exists, we will also point out and emphasize that change is part of that order. We will discuss the science of change, which involves dynamical systems theory, and mention how this area of study will impact our task. We will also discuss how the field of artificial intelligence has attempted to automate intelligence, and we will look at some major developments and promising directions within that complex field.

One set of questions to ask of our model would be "What is the product or output of the model? What happens to the input? What drives the process? and Where is the process going?" These and other questions have been in the back of our minds throughout our work. It is interesting to bring them to the forefront and examine them. We will address these and other questions in this report. Some of the other things we will address are how we think and how we learn. The very nature of intelligent information processing will be examined in light of our task at hand. The reasons for examining the nature of intelligence are to come up with useful modeling ideas as well as to guide our efforts. In our attempts to automate intelligence, the best example for us is somewhere between our own ears.

As part of our discussion we will talk about what we call the nature of data, an obscure topic with much potential in the information processing domain. We will also introduce the data spectrum, which describes how fundamental components of data are structured and related to each other, and the octave rule, which has to do with ranges and order present in the natural processing of data. The input to our model or network is generically called data, but for clarity we will provide a better definition. We will also provide definitions for other important terms, with the realization that trying to define abstract terms like intelligence and information may be counterproductive.

Once we have discussed our approach to modeling intelligent information processing, we will address how we may develop and direct this approach to solve problems in specific domains. The underlying processes of intelligent information processing no doubt affect all disciplines, and efforts to automate the processes will impact each discipline in a different manner. The difficult part is discovering the basics. Here we are particularly concerned with the disciplines of reliability and electromagnetics. We will see how looking at the fundamentals of reliability, electromagnetics and information processing will be mutually beneficial. We believe that in time interest and activity will increase across these areas, with much overall benefit to the technological mission of the Air Force.

Work involving the automation of intelligent information processing has significant long term implications. The Air Force cannot ignore the potential that the so-called "intelligent technologies" may have on its mission. At Rome Laboratory, much of our work revolves around computer simulation and modeling. We hope to develop more efficient ways to process data, in all its forms. No matter if our concern is reliability or electromagnetics or any other relevant technology area, we will constantly be challenged to address "intelligent technologies" in general and to develop specific ways in which the Air Force will benefit from them.

1.1 Order, Organization and Information

We have already mentioned that structure and order exist in our world, and we hope to exploit them in our attempts to automate intelligent information processing. This structure and order is evident everywhere, given the proper perspective. Order exists external to the brain, and it also exists inside the brain. The internal order of our brain is actually a model which represents the external order we perceive. We achieve a natural harmony, an understanding, when the two line up, so to speak. The internal representation in our brain is made up of matter and energy, just as everything else is. But there is more to the story. There is something else which results from the dynamic actions and interactions of matter and energy inside the brain. We generically call this product information, but there are other names for it also.

We understand that the entire universe consists of matter and energy. Let us now assume, for argument's sake, that *information* is also a basic property of the universe. Information is used to describe such things as distance, time, motion, direction, etc., entities which cannot be defined in terms of matter and energy alone. Information is an additional requirement which describes and relates the order and organization present in all forms of matter and energy. Information is a kind of language for the universe. The word, taken literally, implies the reduction of uncertainty. If information actually does reduce uncertainty, then the term will serve our needs quite well here.

The issue of information being of similar status in the universe as matter and energy has been virtually avoided until now (Stonier, 1990; Wiener, 1967). It has taken the advent of the modern computer and man's attempts to build intelligence into the computer to make us realize that this issue requires further attention. The distinguished scientist and scholar Norbert Wiener said that the job of the scientist is to discover and explain the order and organization present in the universe (Wiener, 1967). If we want to somehow model that order, we must first acknowledge that the order exists, attempt to describe its nature, and then capture the essence of that order in a model. This would inevitably involve a process, and the process would have to be dynamic. In this case the process transforms what we call data into something more useful, i.e., knowledge. As humans we interact with this process daily, and we can even change its course. What this

really means is that we are changing the order in our model, our internal representation of the world -- we are changing our mind! We say that the product of this process is knowledge. Actually, we could use other terms like information, wisdom or truth, but the point is that the initial raw material, the data, gets converted into something more useful. The set of processes which allows us to achieve this involves transformations, and we acknowledge these transformations in light of the ever present order, organization and information around us.

We must mention that our usage of the term *information* is very different from the context of Shannon (1949) as well as that of communication and information theories. In their context, information concerns the nature of the transmission of signals in the presence of noise, over a communication link, without regard to the inherent meaning of those signals. In this report we are not so much concerned with communication and the actual transmission of signals, but with the content and especially the meaning of those signals with respect to a process which leads to understanding and intelligence.

1.2 Change

We have mentioned that order and organization are inherent properties of life. Everything goes through change. Some of the change is noticeable, and some is not. (This is mainly due to issues of scale -- time, size, form, etc.). What this means is that things go from some kind of order, through change, to some kind of new order. If you start to consider where or when one kind of order ends and another one begins, it becomes impossible to draw the line. This is because the change, the actual transition, is part of the basic order or scheme of things. One of our tasks is to consider what or where this change leads to. The whole process of change implies a dynamical system, which we will discuss further in Chapter 4. The important point to remember is that change is part of a natural process. We will attempt to exploit the process of change by examining the natural transitions and transformations which occur with respect to data and intelligent information processing.

Change which occurs in nature is often due to an iterative process. This means that the process occurs over and over again. For this to happen, the conditions have to be just right. That is, the root cause of the change has to remain present and the environment has to continue supporting the change in one form or another. As it turns out, this happens often in the brain. Note that in the brain the root cause is not the data or entering signals. The root cause is some biological process. Data is a kind of fuel for the process. When we think and learn, the conditions are usually just right to sustain the many iterative processes and transformations within. However, with change being such an important factor, we can imagine that the current state of our mind is relative. This means that the "state of mind" can and does change, depending on the many forces and influences present at any given time. It's no wonder that we have so many clichés when it comes to describing the mind's processes -- we don't really know what's going on inside! We do know one thing though, that change is an integral part of the process.

1.3 The Nature of Data

The nature of data includes the physical nature of data as well as the resulting cognitive products, the signals and patterns that comprise our mental associations. The concept of the nature of data includes the data spectrum and the octave rule, each of which will be described in Chapter 2, and the harmony of data, which deals with how signal components combine and interact with each other as part of a constructive arrangement or pleasing integration, contributing to the control and order of a more

global system. We will use these and other concepts to show how important order is in information processing and how it may be used to model intelligence.

A study of the nature of data begins with one or more signals. These signals are the result of a disturbance in some physical entity. For intelligent information processing, the signals are either generated external to our body and received by our senses, or they are generated within our body due to some sub-process and are then transmitted as signals within. Either way, for all practical purposes, we can consider them as signals when they reach the brain. This is important because it provides a common ground on which to stand for the purposes of intelligent information processing. We now have a generic starting point that covers all possible sources of information, and we can concentrate on subsequent processes. We need only to worry about the nature of signal interaction. Even though this may not be the actual case in our brain, it is sufficient to get us started. It is also interesting to note that once a disturbance which generates a signal is gone, the only thing that remains in the brain may be that same ethereal entity we referred to earlier and generically termed information. The signals go in and they do something. What is happening? Something is forming.

As signals enter and travel within the brain, they are transformed and transitioned in many different ways. The nature of the network and the forces and influences present manipulate the basic components of signals. These basic components are frequency, phase and amplitude, the building blocks of the signal. Regardless of where they originate, in order to play the brain's game of information processing, all signals (their components) have to conform to the brain's rules. Along the same lines, if we are to build a model which performs intelligent information processing, we must design our system or network to manipulate basic signal components. We have to define the rules of the game. Our model, its very framework, must conform to fundamental principles of nature. By examining the basic functionality of the brain, we can begin to model and develop some of these functions simply by manipulating signal components.

The kind of manipulation of signal components we are referring to gets to be quite complex, as we shall see. After all, we are talking about the man-made equivalent of intelligence. Among other things, our model must be able to handle many different kinds of signals which could interact and interfere with each other in many different ways. We will emphasize the importance of *signal resonance* in our model. Resonance is due to the reinforcement of normal or fundamental modes of vibration. With respect to signal interaction, resonance is characterized by the intense or enriched response to signals (forces) present. The natural vibrations of signal resonance, along with their possible variations, are an essential part of (our impression of) natural intelligence. Resonance may provide a kind of reference for intelligent information processing. We will also use other signal interaction properties, such as interference, harmony and associations, as part of our model.

In a practical sense, the usefulness of such signal manipulation within a model is all a function of design, and we are not suggesting that our model can perform the kind of functionality that the brain can. We are not trying to duplicate exactly what goes on in the brain, but instead we are after only a small subset of the basic processes therein. In our model, we must be able to build up complex signals by combining many simpler components. We will also have to break complex signals down into their fundamental constituents, being careful not to destroy or distort any inherent information. All of this will be for the purpose of getting useful information out of the signals which enter and are transmitted within our model. In the long term, by carefully manipulating the phase, frequency and amplitude components of signals, we believe we will eventually be able to engineer into our computers something akin to the man-made equivalent of intelligence.

2. The Data Spectrum and the Octave Rule

2.1 Describing an Elusive Process

As part of our analysis of intelligent information processing, we initially tried to imagine what kind of processes or mechanisms had to take place at a high level. We knew that signals entered the brain via the senses, and considering sight and sound, these signals were oscillating at high frequencies. We say high with respect to those of thinking or communicating, which are orders of magnitude slower in frequency than the signals constituting what we see or hear. This indicated that a set of reductions or transformations must take place. Some kind of process, or set of processes, was taking signals with high frequencies and converting them into new signals with inherently lower frequencies. Upon further consideration, we developed a framework or conceptual structure to portray this process. This framework represented the various forms of these signals as well as their inherent relationships. We call this framework the data spectrum.

The data spectrum provides a unique perspective from which to view intelligent information processing. As part of the concept, incoming high frequency signals are converted into something with lower frequency, but more importantly, something with more value. Although the data spectrum provides us with a useful framework, it does not fully describe the process. Subsequent consideration revealed another related process which allows the many signal reductions and transformations to take place in an orderly fashion. We call this process the octave rule. Together, the data spectrum and the octave rule have helped us to better envision and describe an elusive process, namely -- intelligent information processing.

2.2 The Data Spectrum

Intelligent information processing is a dynamic process. It ultimately involves the transformation of signals into something meaningful and useful. In an attempt to avoid confusion, we will begin by defining some important terms. Data is the term given to the most basic class of signals. Data is input, raw material or fuel for the entire process. Any class of signals with frequency, phase and amplitude components can generically be called data. Information is a term which describes an orderly arrangement of data, and thus is a more organized form of (the generic term) data. Knowledge is defined as organized information. With knowledge comes more refined associations. Knowledge is that which relates data or information to something of more value, such as a particular meaning or an inherent association. Intelligence is the ability to acquire knowledge and to effectively apply that knowledge in a changing environment. All of these concepts hinge upon the fact that during intelligent information processing data is transformed into knowledge.

This forms the basis of the data spectrum. As shown in Figure 2.1a, raw data signals get combined and organized into information, and then may get transformed into more complex associations called knowledge. The data spectrum portrays this transition of signals from some initial form to some final form. But in reality, the form is relative, because at any given time, all that matters are the associations which exist with respect to the signals present, in whatever form, at any given time. In a larger sense, what really matters is where those associations are going, or what they produce. This dynamic process involves learning and produces knowledge, enabling intelligence (Figure 2.1b). With respect to the overall process, we may assume that the product or output involves communication. However, as we have implied, the real product of intelligent information




Figure 2.1 (a) Basic data spectrum shows data elements as a function of frequency.
 (b) Transformations of data signals enable intelligence.

processing is not communication, but each and every complex association which ultimately comprises intelligence.

We were initially interested in a specific aspect of intelligent information processing, namely, the aspect which concerns the frequency component. We knew that frequency had to play a key role in the overall process. We felt comfortable with the natural implications of this since time and frequency are closely related concepts. As is

portrayed by the data spectrum, time-sensitive data gets transformed into less time-sensitive information, and then into even less time-sensitive knowledge. This means that the importance of time is somehow being reduced (or concealed). The resulting associations are becoming more generic, more independent of time. Herein lies one of the strong-points of the process portrayed by the data spectrum -- the process transforms but does not remove the element of time (or frequency) from the signals. Time is built into all forms of data. The transformed signals are less time-sensitive, more general. But they still have the essence of time built into their complex associations.

This concept of generic associations is very important with respect to the data spectrum, because it leads to generalization. In the process of transforming many signals with relatively high frequencies, the data spectrum implies that the process produces fewer signals with lower frequencies. But perhaps most important, the resulting signals have something which did not exist in any of the original signals. The newly formed signals have inherent associations, general concepts, and main points built right into them! One of the essential features of intelligent information processing is to allow us to make sense out of the volumes of so-called information we deal with daily. The natural tendency to generalize in an orderly fashion is reflected in the data spectrum.

We can see that the word data has already taken on multiple meanings here. One definition for data is the most basic form of input, as in raw data. Another comes from using data as the root word of a more general concept which includes many forms, as in the data spectrum. Unfortunately, we will be using both meanings in this report. As for order and organization and the data spectrum, we have seen how the transformation of raw signals into more orderly forms are portrayed by the data spectrum. The price to pay for this order is the amount of control and effort needed to accomplish meaningful transformations. We can envision a model of intelligent information processing having specific features by design. One such feature is an architecture with layers, each performing some subset of the overall process. In going from signals with high frequency to those with lower frequency, we naturally envision a filtering process which keeps track of these so-called frequency ranges. The filtering process has to be able to focus signals into natural ranges for subsequent processing. This requires an orderly process which we call the octave rule.

2.3 The Octave Rule

The octave rule is part of a natural process which controls the interaction of signal components in intelligent information processing. The octave rule helps scale, filter, focus and group data into specific intervals. The most common interval or ratio is two to one. The octave rule is a generalization of the musical concept of an octave, where a group or range of signals exists within a ratio of two to one (that is, the largest frequency component in an octave is twice the smallest). The important point is that a certain kind of order exists, and this order is not just found in music. We believe the octave rule is part of a more generic process. The general concept of an octave can involve any ratio, not just two to one. Also, the general concept is not limited to frequency components. For example, in image processing, amplitude or size seems to conform to the same kind of principle. This has to do with the importance of scale with respect to the modeling process. We will talk more about the importance of scale later.

The octave rule is part of the orderly process of reducing many signals to fewer, more general signals. Thus it is a scaling process. Control of the process causes signals to form into natural groups or ranges, allowing for more focused processing. The octave rule is ultimately part of an iterative process which involves feedback to help control the

process to stay within an octave. We envision it as part of an internal system of checks and balances which helps maintain order. The octave rule has allowed us to conceptually break a very complex process down in an orderly fashion. The process of interference involving many signal components within an octave results in fewer signals. Although there are fewer signals, each one has a richer, more complex set of associations. More importantly, we believe that the similarities and differences between the resulting signal associations within an octave play a key role in intelligent information processing. We shall see that within an octave, associations exist which form the equivalent of reference and range, which leads to what we call "concept within a context," described more in Chapters 5 and 6.

The filtering and focusing processes which occur at various levels in the brain result in relatively few output signals as compared to the large amount of signal activity going on in the brain at any one time. By virtue of an octave involving a relatively small ratio, say two to one, signals within one octave are already close to each other, so to speak. We believe this plays a major role in intelligent information processing because it allows the exploitation of signal similarities and differences within a close range. For instance, by helping reduce, filter and focus many signals into increasingly finer ranges, the octave rule can help highlight relative differences of remaining signal components. These signal differences can be used to further process and more efficiently control necessary signal interactions. Signal similarities can be used for reinforcement. The overall process, however, must be such that significant signal components are maintained and insignificant ones are eliminated. This will be a difficult task for our model to perform. But for humans, this kind of signal interaction has something to do with conscious decision-making, along with unconscious activity as well. Whatever the case, all this happens naturally, mainly by virtue of the processes of constructive and destructive interference. All in all, the concepts of the data spectrum and the octave rule can help describe intelligent information processing, where data gets transformed into information and perhaps into knowledge in the complex processes of intelligence.

3. Artificial Intelligence and Neural Networks

The next part of our discussion involves artificial intelligence, or AI, which may also stand for automating intelligence. The discipline of AI has been around for almost forty years, with a lot of activity and many developments. However, as the history of AI has shown, the road has been long, with many ups and downs. Controversy has existed right from the start concerning which approach to take, what constitutes success and even whether or not it is possible to automate intelligence. One such controversy concerns the discipline of neural networks. For example, we believe that neural networks is a sub-field of AI. We say this because both areas are trying to do virtually the same thing, using different approaches and working at different levels of abstraction. Both areas attempt to model certain aspects of human intelligence using computational methods. The more generic term for this is artificial intelligence. However, we do not wish to create terminology problems. We shall see that while the approaches of these two disciplines are quite different, benefits will be gained from work in both areas.

3.1 *Attempts to Automate Intelligence*

At a high level, in order to be considered intelligent, a model must be able to perform certain tasks, such as inputting, outputting, using, storing and learning information in an effective way. There are many ways one could approach the development of such a model, and certainly many have been tried. A good account on the history of AI is given by Crevier (1993), and Arbib (1995) provides an interesting view of past and present work in many disciplines related to neural networks. Our work, however, has been influenced by a general sense that something very important was missing from the big picture. We have struggled to find a proper foundation for our work. In the very least, we look for an appropriate direction in which to proceed. This report reflects our general approach and perspective on how to automate intelligence.

We have defined intelligence as the ability to acquire knowledge and to effectively apply that knowledge in a changing environment. This implies the existence of goals and a process for achieving these goals. As for a changing environment, AI codes have traditionally had much difficulty dealing with change (i.e., learning). Unfortunately, nothing in AI or neural networks is as generally applicable or widely accepted as we would like. Today's programs are nowhere near being able to do what the brain does naturally. An even more challenging issue concerns the full potential of the brain. While we believe human abilities are far from being fully tapped, lagging quite far behind is the automation of some of the normally recognized abilities.

The main approach taken by those in conventional AI has emphasized symbols, or formal symbol manipulation, in attempts to automate intelligence. Out of this approach has come many developments in such different areas as logic, reasoning, planning, knowledge engineering, representation, recognition, control, language and rules. Research in neural networks, on the other hand, has concentrated on basic architectural considerations including such areas as weighted connections and layers of processing elements. The biggest promise of neural networks is the ability to incorporate learning, mainly by virtue of their adjustable weights and neuronal transfer functions. While both conventional AI and neural networks have very different approaches, both disciplines work toward a common goal -- to embed aspects of intelligence in computers. The view that conventional AI and neural networks comprise different parts of a single big picture is illustrated in Figure 3.1. This illustration is an extension of the data spectrum concept.

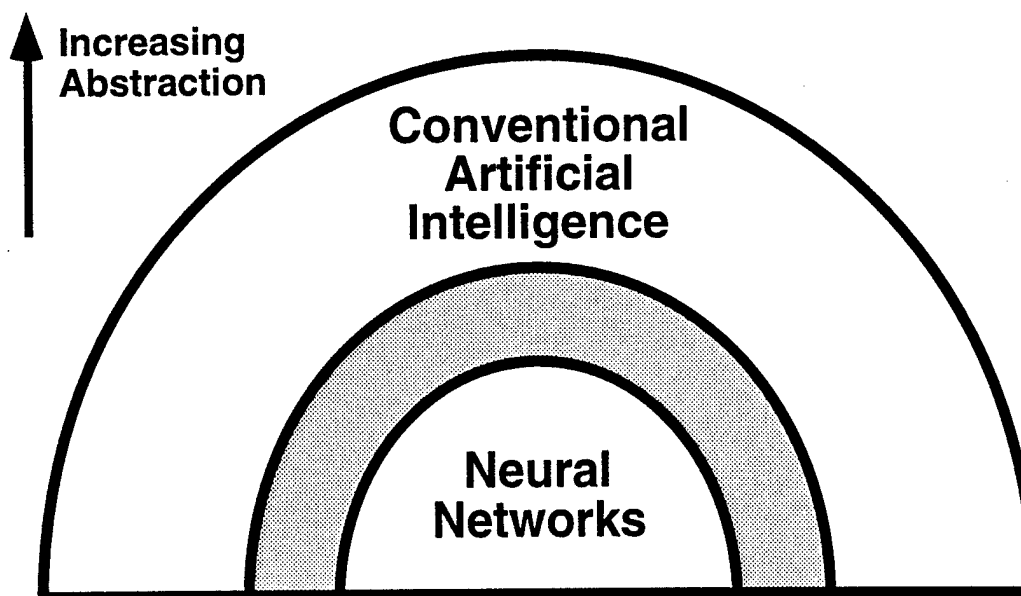


Figure 3.1 Hierarchy of artificial intelligence.

The normally perceived functions of intelligence reside in the outside layer of Figure 3.1. We consider external communication and language to reside in this region. This includes all forms of communication, regardless of which of the senses is involved. It is at this highest level that we are able to manipulate symbols, examine cause and effect, and in general to perform actions normally attributed to conscious thought. However, underneath this functionality there must be a framework or supporting structure. The very essence of this framework should support the many relationships among signals, the so-called associations, used in intelligent information processing. Together, the internal structure and workings (e.g., neural networks) and the external layers and functionality (e.g., conventional AI) work to produce intelligence. Emphasis in conventional AI has been mainly at the external level, while that of neural networks has been more at the internal level. Each approach has failed to capture the essence of intelligence, yet both contribute to the task at hand. We feel that both approaches are needed. However, even taken together, there are still important pieces missing. At the very least, links are needed between conventional AI and neural networks which will add functionality to each approach as well as enhance existing features.

3.2 Lessons Learned

Much has been learned over the years by researchers in both conventional AI and neural networks. An understatement is that the task is difficult, and good results are hard to come by. In a technical area that doesn't have a solid enough foundation, we lack sufficient theory to give a detailed explanation of intelligent behavior. Consequently, attempts to automate or simulate intelligence have not done too well. We certainly cannot declare success, if success is to be measured by widespread acceptance. Of course, there have been many successes on a smaller scale, depending on who you ask and what their definition of success is. But all in all, we still need a better foundation from which to build our applications.

One major criticism of AI programs has been for working on problems in very limited domains. The problem stems from the nature of intelligence, and its dependence on understanding the meaning of (many diverse) important concepts within appropriate

contexts. As a result, researchers have had to severely limit their problem domains in order to progress. Another thing we have learned is that, as far as intelligence is concerned, a *process* is at work. This process involves the transformation of signals in various forms into something very meaningful and useful (i.e., associations). We know that different approaches to AI concentrate on different aspects of intelligence. We need to search for common or primitive elements of intelligence from which we can develop better models. We should integrate useful portions of existing theories and also fill in for missing functionality by proposing and developing new ideas.

One of the most important aspects of intelligence involves associations. We envision intelligence as consisting of many kinds of associations. Conceptually, associations relate inputs to outputs. From a neural perspective, associations result from many inputs being combined by neurons, which then produce outputs that are some function of the inputs. The realization and implementation of associations in engineering models is the subject of much research. As the brain develops over time, we envision complex associations building up and being embodied in the brain's network. Together the brain's architecture and operation enable it to learn just about any kind of association. Associations enable the brain to organize, recognize, and analyze data. It can classify, optimize, estimate and evaluate. The brain can also generalize, memorize, predict, model and control. Somehow all of this results from basic signal components interacting with other network components. We have mentioned earlier that signals entering and traveling in the brain contain information in various forms or stages of development. This so-called information exists as signal associations. Associations (input-output mappings) exist at and across various levels of the brain, resulting from combinations and interactions of signal components as they pass through and between neurons. At a higher level, associations build concepts. From concepts come all sorts of intellectual baggage. We believe that the stringing together of associations and concepts leads to consciousness and thought. Without the underlying mechanisms involving signal associations, it would appear impossible to attain the higher level constructs of intelligence.

We now briefly mention a few lessons learned from some of the many different areas of conventional AI. Various automated methods which make inference or draw conclusions have been developed, but new ones must be able to identify main points of a concept within an appropriate context. This requires mechanisms for reference and range, as well as ways to address the ever-important issue of meaning. As for mechanisms which make decisions, they must be made to do so at the proper level or scale, which depends upon the situation, or signals present. We believe that a process-oriented approach will help in the development of better models. As for rule based systems, which have been made useful in many applications, these systems must eventually be augmented with the ability to learn, or at least the ability to handle exceptions and variability. Learning can help compensate for rigid rules. As for logic, various forms of formal logic have been developed over the years, but for the most part they are limited by their lack of emphasis on meaning. Various methods to approximate reasoning as well as methods which address meaning of a concept within an appropriate context will help. Knowledge acquisition is another area of concern in AI, with sources of knowledge and methods of representation important considerations. One way to address these concerns is to build knowledge bases from the bottom up, using basic elements or knowledge primitives and common features (e.g., statistical methods).

Many lessons can be learned from work in neural networks also. Various methods of learning have been suggested, with much emphasis on supervised forms of learning. However, unsupervised methods are also needed to handle learning in a more natural and unrestricted way. Networks which allow this are sometimes referred to as self-organizing systems. The most difficult networks to design may turn out to be the

most robust and versatile. For example, the functionality of the brain involves nonlinear transfer functions, feedback, parallelism and asynchronous operation. How can we model this? We will eventually have to incorporate dynamical systems concepts to model this kind of functionality. As for binary-based models and programs built to run on digital computers, we must develop methods to better handle approximations (e.g., fuzzy logic), uncertainty and probabilities, and not just the extreme cases of "0" and "1". With respect to the decision-making process, we would like to mention one aspect of the serial/parallel processing issue. Regardless of what processes may be performed in parallel, the overall decision process may be modeled as a single decision mechanism which occurs in a serial fashion. Such appears to be the case in human decision making, with people making conscious decisions one at a time. We may, to some extent, perform input, output and background processes in parallel (Pashler, 1993), but when it comes to conscious decision making, we seem to do this serially. As for timing considerations, synchronous is precise and orderly, but the world is not synchronous. While everything in the world is affected by time, they are not all lined up in time, so to speak. Our model must be able to handle asynchronous inputs and situations. With respect to architecture, hierarchy is important, as is the ability to scale up and down. This implies the need for a layered architecture and a scaling process (e.g., the octave rule). The neural network model of processing elements connected by modifiable weights seems to offer appropriate mechanisms for representing associations.

We have only mentioned a few general concepts concerning over forty years of research. In time, better models will be developed, and they will have to be tested. Testing involves identifying processes, monitoring and controlling them, knowing their outputs, and understanding their fundamental behavior. Models will be developed, applied, and then forced to evolve. It is often helpful to look at past developments. Also, it can be especially important to consider where things appear to be going. It is in this vein that something very important appears to be missing from the collective work in conventional AI and neural networks. We touch upon this issue throughout this report, and although we do not provide all the answers, we offer a few suggestions, and also try to ask some important questions.

As for the state of the art in computer technology and its impacts on AI, there will always be researchers who say they need more computer power to enable whatever approach they're advocating. However, the amount of power is not as important as the clever and efficient use of that power. The usefulness of any technology is a function of intelligence -- we must make the most out of what we have. We have plenty of processing power right now. What we need are better models. As time passes, computers will become more and more powerful. But we must overcome an enormous stumbling block, and that is, how to represent the basic functionality needed to model fundamental aspects of intelligence.

Over the years, failure to capture the essence of intelligence in machines has generated excuses involving such things as inadequate theories and models, programming languages and compilers, and even the state of the art in hardware and software. Certainly these excuses, or reasons, may be true in very specific situations, but in general we will always be confronted with limitations of some kind. Each approach will necessarily have its own limitations. By the way, there have been many different approaches aimed at automating intelligence. Attempts have come under many names, each having a slightly different goal or emphasis. These include such areas as expert systems, fuzzy logic, approximate reasoning, genetic algorithms, evolutionary programming, artificial life and computational intelligence, to name a few. The bottom line is that we are quite a ways off from truly automating intelligence. We must look for commonality among existing theories and build from it. Artificial intelligence and

related technologies are tools. They are enabling technologies that are still being developed. This is a volatile field, and given the information age, we are in especially trying times. Knowing this, we must handle these information technologies with care.

3.3 Future Directions

We have seen that many approaches exist in the attempts to automate intelligence. We have mentioned only a few here, with others yet to be considered. Our main concern is in developing a useful, versatile model. We try to use natural concepts or processes, i.e., those observed in nature. However, when doing simulation, certainly anything that works should be considered. Design and development of a model should address overall behavior. This includes the behavior of internal components, as well as how these components interact with the external world. In the computer world, we must consider hardware and software issues, as well as the gray areas in-between. We must also consider both short term and long term effects. And last but not least, our model must be able to handle change. Change is natural, but there is order even in change. Order should be exploited in any way possible. Byproducts of order are harmony and resonance, important concepts in our model. Learning and adaptability are also key concepts. Generalization is another fundamental component to higher intelligence, as is the use of complex associations. Better automated use of sensory data is also crucial. Without this, common sense for machines seems impossible.

Our approach to modeling intelligent information processing has emphasized the use of signals rather than the manipulation of symbols. As such we see the frequency, phase and amplitude components of signals as fundamental building blocks for computational models. Signals may be represented and processed in many different ways and forms, both spatially and temporally. Pulses are one way to represent, transmit and process signal components while maintaining both space and time characteristics (Johnson, 1994). Pulses may be used to model the equivalent of action potentials in the brain as described by neurobiologists (see, for example, Kandel and Schwartz, 1985). Coupled oscillators and state attractors, described in the next chapter, may prove to be very beneficial in future computational models. Also, by combining fundamental time and frequency domain characteristics, resulting transformations may allow for more efficient processing (Hlawatsch and Boudreaux-Bartels, 1992). As models become more and more advanced, higher levels of abstraction and representation are needed, such as the use of symbols and objects, to help capture and communicate the essence of intelligence.

In our attempt to automate intelligence, we are not after a computer program, *per se*, but instead we are after a computerized process. The *process* is the key. We are after an orderly process which transforms raw signal activity, or data, into something more meaningful and useful, such as knowledge. As such this process must exploit order. The process can be broken down into sub-processes. Throughout this report we offer our perspective on how to begin modeling these processes and sub-processes. Certainly more work is needed, as are other perspectives. We must learn from the past, and keep an open mind. As we have already mentioned, one of the things we need to do is form better links between conventional AI and neural networks. Neither conventional AI nor neural networks, as we know them today, can do it alone. Each has very desirable aspects. We are all examples that intelligent information processing requires at least these two levels of functionality. Why shouldn't our models?

4. Dynamics and Coupled Oscillators

We now turn our discussion to an area which involves dynamics. All systems exhibit a certain degree of dynamic behavior. Dynamics involves change, which often implies a process, such as iteration. One may model and predict the behavior of many processes using simple iteration. Given the proper function, iteration can produce very intricate and complex order, such as in the Mandelbrot set (see, for example, Gleick 1987 and Penrose 1989). We are interested in modeling dynamic behavior. We would also like to include some form of automated intelligence in our model. Thus we investigate certain dynamic aspects of brain behavior. In the long run, we would like to know what to expect of our system's behavior as it changes over time. This involves the study of dynamical systems. We will necessarily limit our discussion here to certain basic concepts.

Dynamical systems theory allows us to describe the behavior of a system as it undergoes change. This is done by describing the properties of the system under study using a set of parameters. By changing the values of these parameters, the corresponding response of the system will trace out curves, or trajectories, which reflect the asymptotic behavior or limits of these values over time. Whether used in the experimental sense or in the modeling sense, dynamical systems theory can be used to help design, control and predict the behavior of a system as it undergoes change. One of the uses of dynamical systems theory is to predict system behavior using a model, with the behavior of that model adequately reflecting actual system performance.

As a dynamical system undergoes change, its behavior is such that individual parameters approach state attractors on their way toward (a more stable) equilibrium. We may use concepts such as state attractors, state space and phase portrait to help model the behavior of a system over time. State space refers to the set of all possible states of a system. Each state has an associated trajectory which represents a possible path leading to an attractor. An attractor is a stable state, one which the system approaches over time. Together state attractors, trajectories and state space can be used to develop what is called the phase portrait of a system. A phase portrait portrays the total behavior of a system, which includes the overall tendencies of its parameters (e.g., position, motion, direction) over time.

4.1 *The Nature of Changing Data*

A dynamical system is any system that moves or changes over time. Given this general definition, we can imagine that everything is part of some dynamical system and is thus going through some kind of change. Even stationary objects, which have the appearance of not changing, are undergoing change at some level or scale. Thus in order to be as realistic as possible, our model should address change and/or the scaling issue. This requirement is obvious when considering the functionality of the brain. We realize that, depending on what you are trying to accomplish, a more static model may suffice. We have already talked about the nature of data in Chapter 1, which included such concepts as the data spectrum, the octave rule, and system harmony. We know that signals entering a system are subject to the rules of that system. These rules must ultimately reflect natural order and the laws of physics. We have also mentioned that change, by definition, is an important part of intelligence. Thus we search for methods which allow us to represent and describe the changing nature of data.

We have mentioned that the process of iteration is a simple example of a dynamical system. The process begins with a seed or an initial state, and involves some function being iterated. After each iteration, the results may be fed back into the function for the next step of iteration. When dealing with nonlinear functions, simple iteration can breed very complex order. This happens all the time in nature, and is especially obvious in living systems. Each form of life begins with a seed, and through processes involving (seemingly simple) iteration, complex life forms emerge. The point is that iteration is important, but even more important, a *process* is at work!

Our model is based on the premise that as signal components transition from level to level, from state to state and from one form to another form, the overall behavior (process) can be modeled using a computer. Concepts such as the octave rule and the data spectrum can be helpful in describing the transformations and transitions which occur. By examining the nature of changing data, the underlying processes involved may provide clues on how to model certain aspects of intelligence, such as the decision-making process. The word transformation taken literally means to change form. The root word *form* is an indicator of what should be important to us in efforts to model intelligent information processing. The boundary regions which exist around stable states or forms of data may be used to model the decision process in part by considering a decision as a transition or choice between possible stable states. Each state would ultimately have some kind of boundary associated with it. A decision would involve crossing a boundary, or threshold, by varying some set of parameters. This is a very simplistic view, considering that boundaries themselves can change. However, we must come up with better ways to model the decision-making process; since decisions are fundamental to intelligence. All of this involves trying to characterize the nature of changing data and embedding this characterization into a workable model.

4.2 Oscillators

The operation of the brain involves much dynamic behavior, with signals traveling and fluctuating in a very complex fashion. However, there must be an underlying order to these fluctuations and transitions. Oscillators may be a natural way to represent the vibrating and periodic nature of changing data in the brain. Generally speaking, an oscillator is something whose response varies above and below some mean value. Since we are interested in signals with amplitude, frequency and phase components, oscillators satisfy this requirement. Also, oscillators accommodate other natural characteristics of dynamic behavior, such as feedback, stability, iteration and periodicity. Finally, the concept of an oscillator is akin to that of a periodic attractor, since both can be used to model periodic behavior. A stable oscillator is a periodic attractor. Due to the iterative nature of many biological sub-processes, we may use oscillators to help model intelligence. While other kinds of state attractors may be used to model complex system behavior, periodic attractors are a natural first choice.

As a point of clarification, the term oscillator is used here not in the classical (hardware) sense, such as in a pendulum, spring or an electronic amplifier with feedback. We use the term oscillator more in a software or modeling sense, as an abstraction of the hardware concepts just mentioned. In a modeling sense, an oscillator can be used to represent and process dynamic signals and to describe their resulting periodic behavior. In general, signals with amplitude, frequency and phase components may naturally be represented using oscillators since such behavior represents a response which varies above and below a mean value. However, the term oscillator has a connotation of being periodic, which means that its behavior repeats itself. We use the term oscillator in the general or abstract sense, with the simple goal of representing stable periodic behavior.

As such the classical (hardware) concepts of oscillators (spring, pendulum, etc.) may also be modeled in software. The term coupled oscillators refers to the combination or interaction of two or more oscillators. Oscillators, or more appropriately, coupled oscillators, may be used as building blocks in a computational model involving complex behavior. The purpose of this computational model is to describe and predict dynamic (periodic) behavior, with the additional goal of ours to incorporate certain aspects of intelligence into that model.

4.3 State Attractors

We have mentioned several times how important change is, and that virtually everything involves change. It would then be natural to ask "In what direction is everything, or anything, going?" and "How does it get there?" The answers depend on many factors. In general, dynamical systems theory offers us the term *state attractor* to help address these issues. For our purposes, a state attractor refers to a stable state which a system or sub-system tends to approach over time. An attractor is literally that which attracts or pulls together. The physical nature of any system is such that the forces and influences which are present naturally create these so-called state attractors. System behavior tends to approach state attractors, which are a kind of equilibrium state. Here the term equilibrium does not necessarily mean a static or steady state, but implies anything that is relatively stable. For example, systems exhibiting oscillatory behavior, and even certain kinds of chaotic behavior, may be considered to be in equilibrium.

Dynamics refers to the behavior of a system as it moves or transitions between different possible states. The term state attractor is used to describe a general class of stable states. The three most common types of state attractors are the static attractor, the periodic attractor and the chaotic attractor. In theory, the behavior of any system may be modeled (approximated) using state attractors. The first common type of attractor is the static attractor, which is a single point to which trajectories approach. A trajectory is simply a possible path or course leading to an attractor. Once a static attractor is reached, behavior will not change (remains static) unless the system is otherwise influenced. An example of this is a marble placed anywhere in a bowl. Eventually the marble will come to rest at the bottom of the bowl, at a point which can be modeled using a static attractor. The next type of attractor is the periodic attractor. An illustration of this is given in Figure 4.1. The dynamics of any attractor is such that system behavior which starts out sufficiently near the attractor will eventually end up on or near the attractor. In the case shown in Figure 4.1, the attractor is periodic, and is represented by the circle. Any behavior which is stable and repetitive can be represented using a periodic attractor. The behavior of a periodic attractor may, however, be quite complicated before it actually repeats. This figure, though, does illustrate why oscillators are akin to periodic attractors. The third type of attractor is the chaotic attractor, which describes complex behavior which never exactly repeats itself. This is the most complex of the three types of attractors, and generally the most counterintuitive. Chaos means different things to different people, so we will provide a more extensive description of chaos later in this chapter.




Figure 4.1 The circle represents a periodic attractor.

In general, all state attractors describe the local asymptotic behavior of a system (Abraham and Shaw, 1989). Here local refers to the behavior of trajectories on or near a single attractor. Asymptotic refers to trajectories which approach a limit over time. Thus an attractor is a type of limit set. This definition is restricted in the sense that a state attractor receives or attracts most of the trajectories, but not necessarily all of them. Actual behavior is too complicated to describe here. In another abstract sense, of all the types of limit sets possible, only state attractors are directly observable, for instance, by way of experiment or mathematical analysis. As it so happens, many of the concepts of dynamical systems theory are mathematical abstractions which have advanced in conjunction with advancements in computer simulation techniques.

4.4 Modeling Dynamic Behavior

We would like to begin our discussion on modeling dynamic behavior by briefly tying in this subject with our perspective on some of the dynamic aspects of intelligent information processing. As part of normal operation, the brain manipulates the frequency, phase and amplitude components of signals present. One method used to represent the activity of these signals is to use pulses. Pulses may be transmitted, processed and stored in a computational model. Oscillators or periodic attractors may be used at a higher level of abstraction to represent the functionality of many pulses and their resulting complex behavior. At an even higher level, coupled oscillators may be used to represent the behavior of many oscillators as they combine and interact with each other. Our overall perspective involves a hierarchical structure and a multidisciplinary approach to intelligent information processing in which system behavior draws from dynamical systems theory, while architectural considerations draw heavily from neural networks, and higher levels of knowledge representation and understanding draw from conventional AI. No doubt other disciplines will be needed to help improve existing theories and to fill in for missing pieces. The bottom line is that fundamental theories and basic, general ideas are needed to form the common ground on which many disciplines can stand in order to automate intelligence.

Our approach to modeling dynamic behavior has involved many concepts, some of which are quite simple and involve common sense, while others may be much more uncommon and complex. What is really needed, however, is for fundamental concepts to be identified and brought together in a cohesive fashion. This is a difficult task. Existing concepts and ideas must be considered, and new ones proposed. For instance, the idea that waves and linking fields exist with respect to signals oscillating in the brain may provide insight into automating intelligence. Signal pulses have properties well suited for operation with oscillators and state attractors. Pulses may be generated in a periodic

fashion, and the resulting oscillations may create waves or fields of influence. The dynamic nature of these oscillations may be used to model complex yet powerful signal interaction. By modeling this kind of interaction more efficiently we may be able to represent, transmit and process signals more naturally, leading to better techniques for recognizing patterns embedded in data. These waves and fields have certain interesting properties. For instance, neural connections constitute one kind of coupling, but waves and fields introduce another kind of coupling which does not require direct contact. Coupling can occur due to a region of influence or a field of energy which necessarily impacts other model components even though they are not directly connected. Doctors and scientists have long examined the electromagnetic fields and related properties of the brain, but not enough is known about these effects and how they may be modeled at a large scale using a computer. Researchers strive for the equivalent of Maxwell's equations for the brain, or the so-called "Maxwell's equations of thought" (Crevier, 1993).

As part of describing dynamic behavior, we offer a brief description of chaos. The term chaos is difficult to define, so we hope that our summary does more good than harm. Chaos is a naturally occurring phenomenon with quite amazing and surprising characteristics (see, for example, Gleick 1987). Chaos refers to a special kind of order, although disorder may be more apparent. A system can be in equilibrium and in chaos at the same time. Chaotic behavior is characteristically nonlinear, but nonlinear behavior is not always chaotic. In a mathematical sense, chaos is deterministic, but in a practical sense, chaos is not predictable. Chaos is different from randomness, since chaos is deterministic yet unpredictable, and inherently shows signs of complex order. Randomness, on the other hand, implies no order, with the outcome of a single random event being non-deterministic. However, in an average sense, statistical order may exist due to many random events, which can lead to predictability. The order in chaos stems from the very nature of change. There exist physical mechanisms, or laws of nature, which dictate the properties and characteristics of matter and energy (and information) as they transition from one form to another. These transitions or transformations contain order, and chaos is part of that order. The phenomenon is far from being completely understood, but with respect to modeling, the order appears to be a function of scale. By this we mean that, as a process undergoes change, the nature of the variations and the characteristics of change exhibit universal properties of scale. Such scaling properties include spatial as well as temporal aspects. In modeling, scale is relative. In reality, everything we understand is relative, depending on scale or perspective. The universal scaling aspect of chaos stems from the inherent property of having similar ratios among scales, or similar order within transitions across all scales. Whatever the case, order is inherent in chaos. This is very important in the development of engineering models, since we would like to know how our systems will behave over time. Depending upon the circumstances, we should know that our system may unexpectedly exhibit natural chaotic behavior. The main point of all this is that chaos exists, and we must know enough about chaotic behavior if we want to control its existence in our systems. Finally, we have stated that chaotic behavior is unpredictable, and also that with respect to modeling, scale is relative. This is because prediction of a model's behavior depends on, among other things, knowledge of initial conditions. For practical purposes, it is impossible to specify a unique, exact and consistent set of initial conditions associated with chaotic behavior. For similar reasons, it is impossible to precisely predict the behavior of a system when it is operating on or near a chaotic attractor.

The investigation of dynamic behavior seems to reveal more and more mysteries in nature. We use our limited knowledge to overcome our limited knowledge. How can this be? Certainly we are not getting something for nothing. We pay for knowledge through hard work and learning, and we also inherit much from those who have gone

before us. The basic property of a dynamical system is an underlying process which involves change. Intelligence and knowledge acquisition can be modeled as such a process. In our early attempts at doing this, we have learned that there is order in everything, even in change. Order, harmony, the interaction of forces, transformations -- all of these are directed by the laws of nature. Some of this behavior is readily understandable, and some is not. Normal behavior refers to regular behavior, or that which is recognized according to some standard. In the study of complex dynamical systems, we come to realize that there are too few standards. The phenomena are difficult to relate to, and many of the concepts are counterintuitive.

In our early attempts to automate intelligence, we search for underlying processes, examine the nature of change, and imagine the kinds of states that change may ultimately lead to. Dynamical systems theory offers many fundamental concepts which address these concerns. From simple iterations to complex processes involving the combination and interaction of many signals, we can use dynamical systems theory to help us understand the nature of change. The amount and complexity of dynamic behavior contained in our model will change over time, but no doubt dynamical systems theory will contribute much to the modeling of neural behavior. It can be argued that the understanding and modeling of neural behavior is one of the most extraordinary undertakings of mankind. This task will require the efforts of many researchers in seemingly diverse fields of study. Individual goals should be chosen wisely, while collectively the research efforts of many contribute to the task at hand.

5. Intelligent Information Processing

The concept of intelligence can be defined or described using various concepts, each of which may be somewhat difficult to describe. The difficulty in trying to describe intelligence is offset by our belief that these are inevitable steps to take. We believe that in order to automate intelligence it is necessary to consider related processes and sub-processes in order to reveal the basic principles and common functions which underlie intelligence. The concepts we examine here include knowledge, goals, learning, decisions, generalization, common sense, and understanding. We even offer a brief description of our concept of consciousness. Remember that we are not trying to duplicate the brain but instead try to model some of its processes. The main difference is that duplication produces an exact copy, while modeling aims to produce a scaled version of, or useful substitute for, the real thing. Our discussion here is preliminary in nature and offered only as an introduction.

One of the main purposes of intelligent information processing is to establish the truth. One definition of truth is this: truth is that which remains unchanged over time, or stands the test of time. We have already mentioned how data gets transformed into knowledge in our discussion of the data spectrum. Similarly, the quest for truth involves the transformation of time-sensitive data into more and more time-insensitive forms (e.g., knowledge). Along with these transformations come the possibility of variations (e.g., different versions of knowledge). Thus we can imagine our knowledge base as being subject to change or fluctuations. When something seems stable in our mind, it might actually be that it is only relatively stable. Furthermore, when something seems stable or settled, it may be said that it's activation can ring true, or resonate. Resonance here describes the strongest signal activity or dominant vibration modes resulting from forces present at a particular time. We envision this resonance as a possible way to represent the main points of thought, or perhaps concepts. This resonant activity would depend on the interaction of input signals with those signals already stored (e.g., knowledge). As for knowledge being relatively stable (as opposed to being absolutely stable), the dynamics are such that change and variability are built right into the basic mechanisms of the brain. We can and often do change our mind. The most general examples of this kind of change involve learning and forgetting.

5.1 The Quest for Knowledge

We mentioned earlier that one of the main purposes of intelligent information processing is to try and establish the truth. For the purposes of this report, let us say that truth is too elusive, so we will settle for the quest for knowledge as part of our goal here. One definition of knowledge involves understanding what something means. As part of this simple definition, meaning and understanding are crucial concepts. Both meaning and understanding require *associations*. As implied throughout this report, our impression of intelligent information processing and of knowledge in general depend upon this concept of *associations*, which essentially involves relating inputs to outputs. As data is transformed into more and more useful forms, such as information and knowledge, signal components are transformed into associations, which can combine to form concepts. Perception involves activating and recognizing basic associations and related concepts. The associations and concepts stored in the brain represent reality as we perceive it, and enable conscious behavior.

Other definitions of knowledge are as follows. Knowledge is organized information, valuable information, or information that is useful or applicable in a

particular situation. Knowledge enables a determination and an understanding of what really matters. We consider knowledge a foundation for intelligence. As such, the term "knowledge base" seems appropriate. One of the uses of knowledge is to enable the acquisition of more knowledge. This makes the quest for knowledge an iterative process. We consider intelligent information processing a dynamic process which produces something coherent and useful out of a seemingly incoherent mix of signals entering our brain. We can imagine this as a process which tends toward certainty (that is, the process reduces uncertainty). But one might wonder how certainty or precision can ever result from such a dynamic process which feeds off input signals seething with uncertainty and imperfection. One of the underlying purposes of intelligence is to increase knowledge. We have already discussed the importance of order in our environment. We must exploit the inherent order in signals in the treatment of certainty and uncertainty in our model. We must also compensate for noisy or imperfect input signals.

As part of our quest for knowledge, we can gain insight and learn lessons from many disciplines such as artificial intelligence, cognitive science and epistemology. For example, researchers in artificial intelligence have long been interested in developing ways to represent, process and communicate various forms of knowledge to allow for computer-based reasoning. Many researchers have emphasized formal principles of reasoning and logic to better understand and encode the essence of knowledge in their programs. However, a lesson can be learned from formal methods, and that is not to place too much emphasis on form alone. Formal methods certainly have their place in math and science, but with respect to intelligent information processing, formal methods have not sufficiently supported the treatment of meaning along with the importance of form. The realization of what a concept means in an appropriate context is a crucial aspect of intelligence. We need to develop automated methods which can identify main points (of communication) with respect to intelligent information processing, and also be able to interpret the meaning of those main points in appropriate contexts.

From a neurobiological point of view, the quest for knowledge leads to different aspects of intelligence, for example neural development (Kandel and Schwartz, 1985). Normal neural development assumes a healthy brain working inside a healthy body. However, not so obvious but perhaps just as important as hardware is the software aspect of normal development, starting with signals entering the brain. By virtue of environment, neural input or stimulation is defined, and is widely believed to be an important part of neural development. What the brain does with these input signals is the brain's business, but this business has a direct impact on development. We believe that the amount of neural stimulation, and the nature of it as well, are a function of the brain's developmental stage. This means that, for normal development, certain kinds of input are more appropriate at certain stages than they are at others. For example, simple repetition seems to play a key role in early neural development. Later in life, however, the affects of repetition seem somewhat diminished, and sometimes even disturbing. Perhaps this is due to the discriminating qualities which come with development and higher levels of consciousness.

5.2 Consciousness

Consciousness is a contemporary product of capable minds. It is a function of gray matter, so to speak. Unlike knowledge or other forms of intelligence, consciousness is transient in nature. It is not passed on from generation to generation, as far as we can tell. Consciousness is what matters now, and is a characteristic property of the living. This property involves intelligence, thinking, awareness, attention, perception, observation, concentration and communication with respect to self, among other things,

and how one's self interacts with the environment. We do not anticipate or expect machines to ever have the kind of consciousness we humans experience. However, in time we do expect certain aspects of consciousness and intelligence to exist in computers. Time will tell which ones are needed, and which ones are even possible.

In Chapter 4 we implied that an oscillation is a form of resonance among dynamic signals. Considering the brain and intelligence, the environment may be such that oscillations and resonance are natural ways to represent and process changing data. Concepts such as resonance and harmony may be useful in helping us to better understand intelligence and even consciousness. Dynamic interaction and interference of signals using resonant modes and oscillations may lead to possible ways to describe and model some of the mechanisms of intelligent information processing. Again, it is not our goal to duplicate the real thing, but only to model portions of it. We may use the trajectory of a so-called resonant concept as a way to describe and model dynamic intelligent behavior. The trajectory may provide a direction, or a feasible path, and a phase portrait may show many possible paths, and perhaps lead to something equivalent to a "stream of consciousness."

There is much mystery concerning the mind, intelligence and consciousness. These are admittedly sensitive subjects, especially when considered in light of computer modeling and simulation. We do not want to mislead you to think that we propose imbedding consciousness in our model. We examine consciousness to gain insight into some of the mechanisms of intelligence. In order to better understand these mechanisms, we consider such things as order and organization, information and data, transitions and transformation, dynamics and consciousness. All of these must be considered when trying to automate intelligence. By doing this we may develop useful concepts and ideas for better models. We do not ever expect to completely understand the brain's mechanisms, but we may be able to model important aspects. If we can imitate any useful aspects of consciousness, in terms of modeling and simulation, then we will be doing well.

With those words of caution, we now summarize our impressions of consciousness. Consciousness is a dynamic process in which thoughts are produced from a series of concept activations within the brain. Concepts are groups of associations, or input-output mappings, which form within the neural structure. Associations may be modeled in the form of adjustable weights (synapses) which connect processing elements (neurons) in complex ways. In the brain, input signals get combined, processed and filtered at various regions and layers to produce relatively few outputs compared to the amount of signal activity occurring. These outputs contribute to thought. Consciousness involves complex signal activity, or so-called associations, occurring at the highest levels of the brain. We mean highest in the conceptual or functional sense, not necessarily the highest physical sense or outermost layers (i.e., cerebral cortex). Awareness is a high level function which results from a sense of heightened interest in certain kinds of signal activity. Awareness involves a general increase in associations over a wide range of signals. Attention involves a more specific interest in signal associations at finer levels, thus serving to focus awareness.

Consciousness involves many signals coming from and interacting in various regions of the brain. Signals get combined and processed based on previous neural development and on forces present. The development of the brain is very dynamic, leading to the fact that each brain, while similar in some ways, is a unique system of complex processing power. The brain's processes are considered to consist of many sub-processes. For example, the brain must be able to reduce many signals down to a few to be efficient. One such sub-process, the decision-making process, is able to produce a

single response based on whatever signal activity is present within a certain level of detail (i.e., within a certain scale). The need for a single response or decision is similar to why we call on a single judge or arbiter to resolve important issues. A single decision authority is literally the simplest, and often the most efficient, way to resolve an issue. Having more than one judge adds complexity, ultimately requiring a single arbitration process to resolve any remaining discrepancies. Perhaps the existence of such an arbitration process in the brain is balanced by its inherent generalization abilities. The brain represents the ultimate tradeoff machine since it can take many specific inputs and generate general outputs, while it can also take general inputs and produce specific outputs. This ability to trade off between general and specific concepts, and to reduce many inputs to a few outputs in an orderly fashion, is no simple feat. Consciousness depends on these processes.

Another aspect of consciousness has to do with energy considerations. Conscious activity is believed to consume a large amount of energy relative to the total amount of energy consumed by the brain. Consciousness can be described as heightened mental activity. While the body must maintain the overall harmony and balance of very many systems or processes, the brain must control its sub-processes. One such sub-process is consciousness. The location of heightened mental activity in the brain is believed to move around in the neural structure. At any one time, the extent of total activity defines the current "state of mind," at least in a physical sense. Consciousness may involve the highest levels of mental activity (e.g., sequential resonance), but we do acknowledge that much activity is also going on in the background, so to speak.

The brain is able to process many signals at one time. Incoming signals and those stored in the brain can interact in many ways. In order to "make sense" of them all, the brain uses efficient processes to resolve all the signal interactions. Consciousness is one result of these processes occurring in normal waking moments. Consciousness necessarily involves the focusing of energy to produce associations and concepts within the brain. Cognitive development involves controlling which concepts get activated and which associations get stored over time (i.e., where the energy is to be "focused"). Mental energy shifts among the many regions of the brain, depending upon neural structure and signal activity. Output patterns are associations of input patterns. These output patterns trigger other thought patterns, allowing the process to continue in an iterative fashion. Due to physical restrictions, the overall process must allow only a relatively few regions of the brain to be "activated" at any one time. These relatively few regions output even fewer signals. This focusing and filtering is an essential part of the process. The "strongest" signals remaining (i.e., those resonating at this high level) at any given time are believed to constitute or contribute to conscious thought.

5.3 Goals

A goal is a desired end state, one which usually takes much time and effort to achieve. With respect to intelligence, mechanisms must exist which enable the creation and achievement of goals. Once a goal is established, methods must exist which enable us to work toward that goal in an orderly fashion. What could this be? How can we model the concept of a goal, and the processes needed to achieve that goal? Perhaps a goal is a bunch of decisions, and the combination of those particular decisions results in a particular goal whose nature is dictated by the nature of those decisions. Or perhaps a goal is a state attractor, with trajectories serving as paths which lead to that goal. In a physical sense, a state attractor helps organize, control and maintain a balance of forces present (i.e., establish harmony or local equilibrium). Similarly, a goal helps organize thoughts, make decisions and, in general, sort out the prevailing mental activity. This

requires a balancing act of existing forces, along with whatever impact or tendencies these forces may involve. We know that orderly mechanisms must exist in the brain which reduce many signals or priorities down to a smaller, more reasonable amount. The resulting harmony or resonance, which involves goal-oriented behavior, could be thought of as a form of dynamic equilibrium.

Let us consider a little further how goals may be achieved in light of what we know from dynamical systems theory. It is difficult to imagine how goals are physically established or embodied in the brain. But other aspects of goals are also intriguing. For example, for any given goal, how does one come to know what to do to achieve that goal? This is especially mysterious for longer term or difficult goals where no clear path exists regarding how to achieve them. The goal-reaching process must be dynamic. Consider the case where only parts of the process are known, such as performing a task for the first time, solving a riddle, or wanting to go somewhere but having only sparse directions. In each case, you must fill in for missing pieces. In order to do this, you may complete the process by breaking it down into more manageable pieces. Ideally each piece will result in a more understandable situation. In breaking the process down, mental sub-processes are activated, many of which have probably been used before in other similar situations. This implies that certain aspects of these sub-processes are generic and handled in a similar fashion no matter what the circumstances are. For example, appropriate references and ranges must exist to help keep the entire process on track. The dynamical systems concept of a trajectory represents a path of behavior which leads to a stable state, or a so-called state attractor. In a similar sense, a goal can conceptually be considered a state attractor. Also, statistical tools may be helpful in addressing the issues of reference and range associated with a concept. The data spectrum and the octave rule may provide additional ideas on how to model goals.

Another important aspect of goals has to do with the concept of time. Goals are a function of time -- that is, they are affected by time. Depending on what you are trying to do, at any given time, your actions and thoughts are influenced by your goals. Short term goals, as well as long term goals, exist. For instance, expectations may be considered short term goals. What we expect in the short term is an example of how our perception of goals and time come together. The decision process often involves trading off between long and short term aspects of an issue, illustrating the importance of time with respect to goals. The dependence on time also implies that in order to help reach a goal, we may break it down into sub-goals, or smaller increments of time. The important point of all this is that complex processes exist which allow us to create and achieve goals. In order to model these processes, we must come up with appropriate and useful ways to represent the functionality of goals.

5.4 Decisions

One of the main features of intelligent information processing is the ability to make decisions. This ability helps define our purpose, our very being. Decisions drive our actions and help define our goals. While each one of us may not be consistent with all of our decisions, it may very well be that our beliefs and our state of mind hinge upon how well we try. Decisions are required when we must choose among different alternatives. At any given time, our state of mind defines our thoughts, our priorities and our goals. While our conscious activities march onward, the mechanisms which enable us to make decisions work (both in the foreground and the background) to support our needs and our wishes. These mechanisms are so ingrained in our brain that they can and often do operate immediately, automatically. Our ability to make decisions is fundamental to our existence.

Decisions ought to be made at the proper level or scale. One difficulty with this is trying to *decide* what is proper for a given situation. This illustrates an iterative feature built into the decision-making process, with one decision affecting another. As we have mentioned, the issue of scale is critical in the decision-making process. The issue of scale can be applied to just about any concept, making scale a universal property. In a general sense, scale is a proportion, a framework or a relative measure. When we say decisions ought to be made at the proper scale, we are saying a lot. This is because any scaleable factor can affect the outcome of a decision, and all factors are scaleable. For any particular decision, the issue of scale would depend on relevant goals, associations and previously made decisions. Considering the many associations involved in a concept and the cumulative effect of making decisions based on many related concepts, we see that the decision-making process is quite involved.

Since the decision-making process involves scaling mechanisms, there must be links between the different scales or levels involved. Or there must be orderly transition mechanisms which maintain the inherent order and organization present at each level or scale. We envision the octave rule and the data spectrum as aides to understanding such mechanisms. The concept of scale may be used as a modeling construct which helps define the boundary of forces acting upon various forms of knowledge at any given time. The decision process must deal with boundaries and scales, and evaluate appropriate associations and relationships. These mechanisms feed into our ability to draw intelligent conclusions from available knowledge. Relevant factors must be considered, ultimately in the form of positive and negative forces, and tradeoffs must be evaluated. We know that in general these processes are less than perfect. In practice, a lot of variability is built into them. In any event, we try to automate these processes in some form or another. Of course our automated versions will be different from the biological versions, but the important point is that we ought to be able to automate these processes and apply them to useful and important problems.

The decision process involves considering significant aspects of a concept within an appropriate context. That is, many associations which comprise a concept, or many thoughts which come to mind while making a decision, are evaluated with respect to a context or range. Granted, the many associations and their so-called ranges all have dynamic aspects to them. But in general, a concept is decided upon within the bounds of a particular context. We have stated that a concept is composed of many associations. The statistical nature of underlying mechanisms implies that statistics and other branches of math will provide useful modeling approaches. Also, the decision-making process seems to involve mental activity occurring at many levels or different states of consciousness. This implies that we search for scaleable processes, ones which are based on something solid. Often we cannot state our thoughts well enough or explain our decisions, yet we feel confident enough that they are right or well-founded. We search for such an elusive foundation.

5.5 Learning

The ability to learn is one of the most significant aspects of intelligence. Intelligence is a process which tries to make sense out of the world. As the brain senses the many signals entering and being activated within, it tries to map them into its knowledge base. This involves learning, with new associations being formed and existing ones being modified. Order and consistency (harmony) are essential for efficient learning. As mentioned earlier, during the learning process data is transformed into information on its way toward knowledge. The learning process produces new signals

which contain something not present in the original (entering) signals. That is, new signals form which contain significant components or cumulative effects of the input signals. The resulting signals are not merely the sum of the individual inputs. Depending on the situation, the resulting signals may very well be more significant than the simple or linear combination of all the inputs. Whatever the case, the process results in fewer signal components. Whether or not the resulting signals are more significant than the inputs depends on the circumstances. But the bottom line is that the overall process produces signals having significant content or meaning, and these relevant signals can be learned, along with their supporting associations.

In the process of learning, we try to "understand" the meaning of new signals as they become activated in our brain. This understanding ultimately involves signal interaction. New signals (or newly activated signals) interact with existing signals (i.e., the knowledge base). This activation and interaction triggers learning. Whether done consciously or sub-consciously, learning requires an open mind. The mind cannot learn without being open or exposed to new signals (i.e., change). In this context, an open mind means that, for whatever reasons, new ideas may be considered and incorporated into the knowledge base. As for determining the meaning of activated signals, the processes involved are still unknown. We have ideas on how to approach modeling these processes, many of which are mentioned in this report. We believe that one of the keys to this is being able to identify reference concepts within appropriate contexts (ranges). This involves being able to store and recall *relevant* associations based on existing signal activity.

Learning means to gain knowledge or understanding. It also implies the ability to realize cause and effect relationships. Learning may also be defined as lasting change toward improved performance resulting from experience (Yaworsky and Vaccaro, 1993). In any event, learning involves change. With respect to intelligence, that change has a purpose, which is essentially what a goal is. Close ties exist between goals and decisions. As part of the decision-making process, for instance, signals get settled and stored. This results not only in the making of a decision but also in the many associations which support, or are otherwise related to, that decision. The resulting associations constitute that which is learned. Goals, decisions and learning are very closely related sub-processes which support intelligent behavior.

The learning process is used by the brain as it continuously deals with change. Because of the tremendous amount of change present or possible in natural signal activity, the learning process must somehow exploit *differences*. We talked briefly about this earlier in our description of the octave rule. The overall process must not only take into account signal differences but also similarities as well. As signals get filtered and focused into many kinds of classes and categories, as surely must be the case, finer and finer differences become exposed within the bounds of these classes and categories. Together fine scale differences as well as general similarities must be used in the interference process, allowing any level of detail, from general to specific, to be focused on. Within any (general) class, thresholds must exist which rely on signal differences that help distinguish specific instances from general categories. The main point is that basic processes which exploit signal similarities and differences are essential to the learning process. Some of these processes exploit signal differences, from large to small, while others make use of signal similarities.

Reinforcement and consistent behavior are also important factors in learning. Reinforcement can be positive or negative, depending on the circumstances. Consistent behavior is believed to be a contributing factor to harmony. Neural development seems to be most efficient when exposed to orderly and consistent behavior of a particular kind.

An example of this is how we choose what we are interested in. Our interests depend largely on our past experiences. Consistent behavior reinforces associations, becoming part of our experience. These associations and experiences, in turn, tend to reinforce our interests, creating an iterative process. As such, the feedback we receive each time we perform similar behavior provides the reinforcement needed to improve that behavior. Again, the processes are inter-related. By pursuing interests, knowledge accumulates, which contributes to the overall order and organization within the brain.

5.6 Generalization

In the field of neural networks, generalization means that a network can provide an appropriate response to novel inputs. This means that the network can produce an appropriate output given unique or new inputs. The network does this by making associations (input-output mappings) similar to those activated by the novel input. Another definition of generalization is to reduce many associations or concepts down to a few. In doing this, the few (outputs) become the synthesis of the many (inputs). The resulting associations or concepts actually signify trends or tendencies of the input signals. Similarly, generalization can involve taking specific associations or concepts and forming more global or significant ones from them. An overriding feature of generalization is the ability to produce the most appropriate or efficient output for a given set of inputs.

The ability to generalize is an important part of intelligence, one which we tend to take for granted. This ability is powerful, useful and, for the most part, inevitable. In one sense we generalize in order to deal with the vast amount of data entering our brain. We also generalize in order to deal with imperfect data, be it noisy, missing pieces, or otherwise inappropriate. Our ability to generalize is continuously used, and often misused. It is difficult to generalize well. Just as generalization is an important part of intelligent information processing, so is the handling of specific details. We have already mentioned how the brain is the ultimate tradeoff machine when it comes to generalizing and specializing with the various forms of data present. We can memorize and analyze specific signal components, and we can also generalize from them. Exactly what we do must depend on our goals and priorities, and is a complex function of scale. This is one example of how goals, decisions and other mental processes are inter-related.

In theory, one might expect that the products of generalization shouldn't change very much or very often. Generalizations may be modified or refocused, but they should be fairly stable, serving as a foundation for other mental processes, such as dealing with specific details or other transient signal components. Specific details may matter very much at any moment in time, but general concepts or issues tend to matter more in the long run. It is interesting to consider how specific details feed general concepts. By definition the process of generalization proceeds from many specific details to fewer general ideas or concepts. Generalizations are difficult to confirm or support because they usually represent so many different associations, concepts and thoughts. Trying to explain which of these is most important can lead to very subjective results. However, so much goes into the formation of generalizations that they are difficult to ignore. Of course, it helps if the generalizations are appropriate, but they do not exactly come labeled as such. Many times we have difficulty relating to someone because they generalize differently. Each person has a different set of associations and experiences to draw from, and they may even use different mechanisms for generalizing, making communication and understanding difficult at times.

The brain must be able to deal with many timely and specific signals in an appropriate fashion. A natural process must occur which transforms timely signals into less timely, or less time-sensitive, ones. Also, specific signals may need to take on a more general nature in order to be dealt with appropriately. We may memorize or otherwise have to deal with specific instances or items, but at the same time these feed into our generalization mechanisms. Associations which form in the brain tend to take on a general or relative nature. One result of the process is that reality is transformed into a relative version (a model), one whose products and processes depend in part on the current state of mind. One purpose of a model is to *generalize and abstract*. This is difficult to do well. As part of generalizing, signals must be reduced, specific instances must be translated into general concepts, and main points must be identified in order to generate appropriate responses to novel inputs.

Earlier in this report we presented the concept of the data spectrum, which portrays intelligence as a process which transforms data into more useful forms, such as information and knowledge. As part of the spectrum (see Figure 2.1a), the most complex and abstract forms of data exist on the left side of the spectrum, and the more fundamental and basic forms reside at the core. This implies several things. First, high frequency components of data (those changing very often) must be filtered and focused in order to transition to the lower frequency components on the left side of the spectrum. Also, as more and more signals combine to form complex or abstract concepts, some kind of orderly reduction has to take place, resulting in fewer signal components which are more significant. As part of the transformation process, the resulting signal components naturally have lower frequencies than the original components. And finally, in order for signals to make it to the center or core of the spectrum, they have to be basic and fundamental. Due to the tremendous amount of signal interaction which takes place at the core of the spectrum, signal components here must satisfy very many physical constraints. Regardless of their source, all signals here are subject to interference and transformations which result in basic, common signal components. This is an example of how order is naturally exploited in the overall process. The data spectrum portrays the ability to generalize as one moves from right to left in the spectrum, and it also indicates that common sense is approached as one moves toward the center of the spectrum, where many diverse signal components must combine in a common fashion. We will talk more about common sense and how it relates to the ability to generalize shortly.

The brain's ability to generalize may be a scaled up version of what each individual neuron does. That is, the job of each neuron is to take many inputs and combine them in some complex fashion, synthesizing an output which is representative of the particular inputs processed. Similarly, the brain generalizes by taking many inputs and combining them in a complex fashion. Exactly how many inputs, neurons and layers are needed to allow generalization for a given task or problem is not known. But it is interesting to wonder if the ability to generalize is the result of the activity of just a few interconnected neurons. Another interesting note about generalizing is that if we use the wrong inputs or somehow misuse the process, we may very well be misled by the results (our generalizations). It is easy to generalize, but it is much more difficult to generalize correctly. No doubt time, energy and concentration are contributing factors to the generalization process. It is also interesting to consider whether general concepts or specific details are more important as part of the trade-off process. Unfortunately for modeling purposes, the answer must depend on circumstances (i.e., meanings within contexts).

5.7 Common Sense

By definition, common sense is basic consensus or general agreement among two or more people. The implied basis for this consensus is common understanding of sensed signals. Common sense implies that the inherent message or meaning of sensed signals is shared. The message or meaning consists of associations, and common sense is based on common associations built up over time. The associations are usually related to common or natural phenomena, and not highly dependent upon complex or advanced levels of understanding. In order for signals (signal components) to be in common, they typically reflect basic principles, and usually represent generally accepted concepts. Thus common sense refers to *a level of understanding which is (or may become) common among many people.*

Sensory data is very important to our existence, at least from an intelligence point of view. The brain naturally becomes dependent upon sensory forms of data. Once we become used to receiving a certain kind of sensory data, the lack of it may lead to psychological problems (Caudill and Butler, 1990). If a source of data is removed, the brain must adjust or compensate for it, which may lead to much difficulty. A positive aspect of sensory data is that it can be used to help relax an overworked or confused mind. For example, one way to relax is to "come to your senses" by exercising or looking into the distance. This usually helps shift mental activity into a more relaxing state. Finally, our senses are an essential part of our intelligence, enabling such things as communication and perception.

Common sense is an artifact of signal interaction. Common sense implies common signals and common mental processes. The signals have a regularity about them, and the processes involved tend to be basic and natural. Regularly occurring signals may "sink in" and appear to be basic or common. Or signals entering the brain may be in common with those already stored. Both of these situations result in the reinforcement of common associations. The resulting condition is that external signals have something "in common" with internal signals. One of the difficulties within the field of artificial intelligence involves linking external communication and (high level) cognitive processes with lower level processes which must exist in an underlying (neural) network.

In general, common sense is a function of neural development. For the most part, young children do not have what it takes (the means) to experience common sense. This is because they have not yet formed the associations needed to support common sense. In time these associations build up, forming what turns out to be a foundation common to many (similarly experienced) people. Very young children tend to perceive or experience things as being absolute. Their first impressions may very well be memorized, becoming a "matter of fact." They do not easily understand generalizations or abstract concepts. The more they are exposed to them, however, the more they will learn from them. With neural development comes the ability to distinguish unique circumstances from common ones, and the ability to differentiate between specific facts and general concepts.

Communication relies very much on common knowledge between sender and receiver. We do not usually communicate by exchanging the full extent of a message, such as starting with background information, explicitly stating all the main points and then supporting all of them with relevant information. Instead we usually take many short-cuts. We may assume that our audience already knows much of what we are talking about. Or we may offer to a conversation that which we assume others do not yet

know. When we discuss issues that we mostly agree upon, we may merely wish to reinforce our mutual thoughts. In any event, we communicate in order to get our points across, to answer questions, or to question others. With the goal of communication to enable better understanding, or to make our understanding common to others, communication is truly an amazing interaction of signals. We often have high expectations of the communication process in the sense that we expect so much from our audience. All too often we assume that they understand what we are talking about. Sometimes they don't. Sometimes they understand a particular subject more than we do. We have mentioned that there is the real world, and then there is our perception of the world. To each of us, our perception is reality. What we sense, what we think, our attitude, our understanding, our mental images, our consciousness -- all of these are what really matter to us. But we must remember that we are dealing with a model here. Others have a different model. What matters to us often does not matter to others. When models are very different, communication can break down. We ought to be aware of that which may be common to others and that which is particular to our understanding.

Common sense is related to the ability to organize many (specific) signals into more basic (general) concepts. That is, common sense is related to the ability to generalize. But there are differences also. With common sense, concepts are usually basic, natural and simple to understand. The concepts also tend to be common among many people. With generalization, however, concepts may be quite complex, and each person probably generalizes differently. Generalizations are dependent on the mental abilities of a particular individual. The reference used for generalization and common sense is also different. The implied reference in generalization is *self*, with its unique perspective and particular set of associations. On the other hand, the reference in common sense is *many people*. Common sense refers to signals or concepts understood by many. Thus common sense is a form of generalization among many people. As such, in order to be "common" among many people, common sense relies on the existence of simple concepts and the interaction of basic, sensory signals. Only the most basic signal components are likely to be in common among many different people. Generalizations, on the other hand, tend to be more subjective, since they represent an individual's unique collection of relevant associations and concepts.

5.8 Intelligence and Understanding

We have defined intelligence as the ability to acquire knowledge and effectively apply that knowledge in a changing environment. The acquisition of knowledge involves learning, and the application of that knowledge requires goals. Along with gaining knowledge, learning also implies understanding, or the ability to realize the meaning or significance of that which is learned. Understanding is a foundation, something solid upon which thoughts and ideas (signals) can build. The decision-making process is another important component of intelligent behavior. Only through conscious choice and the consideration of relevant trade-off factors can intelligence be realized.

We have mentioned some of the processes of intelligence, such as setting goals, establishing priorities, evaluating tradeoffs, and making decisions using various forms of data. Other processes involve generalizations, logic and common sense. Each process involves learning necessary associations. Each of the processes may work independently, but more than likely they work in conjunction with each other. As part of the natural interaction of sub-processes which takes place, orderly mechanisms must exist which control the transformation, feedback, scaling and communication of data in its many forms. The underlying processes, using signals as the raw material, are headed somewhere. They all produce something. That product is intelligence.

As part of intelligent information processing, understanding involves trying to establish truths (or partial truths). The ultimate goal is to ascertain absolute truth as opposed to apparent truth. In reality, however, understanding is a relative notion which depends on one's perception, intelligence and knowledge of a particular subject or situation. Perception, intelligence and knowledge, however, are the result of imperfect processes. That is, we as humans are not perfect, and it is difficult for us to *come to know the absolute truth*. From a lower level standpoint, understanding involves the "sinking in" or "settling down" of signals (to stable states within a network). These signals produce associations, which comprise concepts and other cognitive products. Understanding involves being able to grasp the significance of concepts and realize their consequences. This realization naturally involves the consideration of cause-and-effect relationships observed by the senses and activated within the mind.

Some of the products of higher forms of intelligence include language, imagination and creativity. Our job of modeling the brain and automating intelligence involves emulating some of these higher level functions. This implies that we should take clues from the architecture and operation of the brain. Many disciplines have done just that, such as artificial intelligence and neural networks. As we develop a better understanding of the processes of intelligence, at the same time we slowly progress toward having machines which can achieve certain levels of understanding and intelligence.

Earlier in this report we used the term *data* to describe input signals. These signals are often imperfect in some way, either due to noise, missing pieces, or some other relative imperfection due to the nature of their source or transmission. We as intelligent beings must try to make sense out of the multitude of signals entering and being processed in our brain. We have to determine *what really matters*. This requires knowledge, which exists in various forms. We must be able to identify significant components of knowledge, which corresponds to relevant and important signal activity. We must be able to know the meaning of this signal activity with respect to some reference or context in order for comprehension and understanding to occur. We know that these processes are dynamic and depend heavily on a knowledge base and the ability to apply that knowledge base in a changing environment to achieve important goals.

We have talked about intelligence in many different ways throughout this report. Obviously there are many ways to describe this complex phenomena. And certainly, further understanding is needed. However, intelligent information processing is crucial to our existence. And more germane to this report is the fact that computers already do so much information processing for us. It is only a matter of time before they do so intelligently, or at least in a more intelligent fashion. Since this kind of computer evolution is inevitable, we feel it is necessary to begin describing fundamental aspects of intelligent information processing. Over time, better descriptions will evolve, leading to a better understanding of the overall process. Many of our initial ideas will turn out to be incorrect or misguided. But we must start with basic principles. In this report we offer ideas on what these principles may involve. We discuss some complex concepts, but more often than not they are merely the combination of many simple ideas.

5.9 The Importance of Time, Scale and Form

Time, scale and form are closely related aspects of a crucial process. The process we are discussing involves intelligence and the intelligent processing of information. This process converts raw signals (data) into something more useful and meaningful, as

described earlier using such concepts as the data spectrum and the nature of data. As a result of the overall process, data is transformed, or changed from one form to another. Obviously form is important. We have also described the octave rule as a filtering or focusing process, one which involves the orderly scaling of signal components as they combine and interact in a complex fashion. Thus scale is important. And we have seen how time is built into all forms of data. The data spectrum portrays a process which involves time-sensitive data, less time-sensitive information, and even less time-sensitive knowledge. The entire transformation process is a function of time. Thus the aspects of time, scale and form are important, and virtually inseparable. They contribute to order. We all know of situations where, depending on what you are doing, everything has to be just right. One example is how time, scale and form all contribute to what we call perspective. A perspective is a point of view, with *the point* being a kind of reference, and *the view* depending upon scale, among other things. Scale can be a proportion, range, dimension or limit. What really matters, though, is that mechanisms exist which create such things as perspective, and no doubt time, scale and form are important aspects of such mechanisms.

In the physical world, the laws of physics (natural laws) rule. In the mental world, so to speak, consciousness rules. We do not understand the natural laws which apply to consciousness. Thus when we say consciousness rules, we imply that the mental processes of consciousness dictate what really matters to us (within the constraints of natural law, of course) at any given time. What we are thinking, what we hold to be true, and our current state of mind constitute our model of the real world. The model is ideally an accurate (scaled) version of the truth. All too often, however, the model is riddled with incorrect or inappropriate information. The result is a major component of what constitutes human error. The point here, however, is that our model (our consciousness) is a function of time, scale and form.

In our chapter on dynamics, we mentioned how order appears to be a function of scale. When a process undergoes change, our understanding of the characteristic properties of change and the inherent nature of variations depend on scale. We have also mentioned that in the physical world, the laws of physics take care of themselves, so to speak. What this means is that, in nature, scale is not so much an issue. It is only in modeling the physical world that scale becomes a critical factor. Modelers must (explicitly) decide which scale or scales are relevant. In an abstract sense, scale can be just about anything, making it relative. But in a practical sense, scale depends on what you are trying to do and on the purpose of your model. Accuracy is important with respect to modeling, and accuracy is a function of scale.

The octave rule was described earlier as a function of scale. The octave rule involves the orderly filtering and focusing of signals into relevant groups and ranges. We believe these orderly processes contribute to awareness and attention at higher levels, and provide other support mechanisms at lower levels. The octave rule helps manage and control the orderly reduction of many signals to fewer, more significant ones. The octave rule must then allow the grouping of remaining signals according to class or type. Groups are made up of signals with similar components. As a result of the focusing and filtering that has already occurred, signals in similar groups have relatively small differences between signal components. This is important because mechanisms which make use of signal differences must rely on magnitude, but at the same time magnitude is a by-product of scale. The lower level mechanisms needed to support intelligence may turn out to be relatively simple. At a higher level, the octave rule and other generic processes contribute to reasoning and other cognitive functions. Many processes contribute to the orderly arrangement of signals across many levels, from cognition down to the more basic functions. Much activity is going on in the brain at any given time,

with short term and long term consequences and effects at many different scales. All of this processing produces *information* in its various forms. The bottom line is that time, scale and form are fundamental aspects of intelligence.

We know that the transformations and processes of the brain are all a function of time, but we still don't know exactly how to incorporate time in our model. We have alluded to the fact that one of the products of intelligent information processing is a time-insensitive version of sensory input. We would like to point out here that the timing mechanisms of the brain may be quite different from the normal concept of time, such as that which is portrayed or perceived by the movement of a clock. That is, the intricacies of the brain may operate on timing principles which we cannot (yet) perceive or comprehend (see, for example, Penrose 1989). This may involve signals whose frequencies are not detectable by our senses, or perhaps it involves signals we can detect but, for whatever reason, do not know how to process (consciously) in an efficient manner. The transformations and processes occurring in the brain are extremely complex, and may involve spatial, temporal and frequency aspects so different from what we are accustomed to that we can hardly describe them. Physicists and mathematicians may provide us with useful transformations, but we will have to better understand the nature of underlying phenomena in order to adequately model their processes.

6. On Automating Intelligent Information Processing

We have discussed various aspects of intelligent information processing in this report. Now we will address how we might *automate* some of these functions. We will not (cannot) provide a detailed description, but instead offer general ideas on how to begin modeling intelligence. One of the main purposes of intelligent information processing is to integrate and coordinate signal activity within a network (brain) in order to harbor intelligence. As such the brain is a machine. The inputs to the machine are signals. These signals are detected by various senses and transmitted to the brain, or they are generated from within the brain itself. The outputs from the machine are signals which somehow constitute intelligence. Exactly how intelligence is produced is not known, but we believe that certain aspects of intelligence may be modeled using various computerized techniques. Very many different approaches have been taken over the years in attempts to automate intelligence. In fact, so much work has been done thus far in so many disciplines that it is virtually impossible to comprehend and appreciate the magnitude of it all. We must struggle to put together significant components of basic theories as we try to develop useful, efficient models.

Our approach to automating intelligence, as described here, may appear relatively simple. Granted, we have provided only a general, introductory description. Complex details involving advanced mathematical techniques and the latest developments in neural-based technologies must also be worked on. But no matter if general concepts or specific details are worked on, a more solid foundation is needed. By concentrating on general principles and searching for simple processes, and especially by considering intelligent information processing a *process*, we have been able to focus on various aspects of the process that would otherwise be difficult to see. As part of human nature, many aspects of intelligence may become second nature to us. That is, underlying processes tend to get ignored, or at least taken for granted. These underlying processes ultimately involve the manipulation of basic signal components and the transformation of these components into more advanced forms, such as thought, language and imagination. We know that we will not be able to understand the entire process, but we believe we will be able to model portions of the process as time goes on. Due to unavoidable limitations, our models will necessarily be subject to change, and forced to evolve. We realize that neither a top-down nor bottom-up approach to automating intelligence will suffice by itself. We must be open to many ideas, as surely many ideas will be needed.

As used here, the term automation refers to a computer carrying out well defined operations as specified by a program. The program is based on some form of logic. This is somewhat ironic because the *program* in biological neurons is not based on logic *per se* but on underlying processes which, when combined, may be used to comprise a logic. The logic on which the human program is based is not known to us, and may never be. What this implies is that modeling approaches which do not accept this fact (*that something fundamental is missing now*) may be doomed. We must try to build our models on something solid. That foundation, that understanding, has eluded us thus far. We do not understand intelligence well enough to be able to program it outright.

Many people believe that we will never be able to automate intelligence. We must be clear about our purpose here. We are not out to automate intelligence *per se*. In fact, if we can model any significant part of intelligence, then we consider that *progress*. Benefits will be gained in many disciplines by the automation of intelligence. On one hand, we may work towards the lofty goal of automating intelligence, knowing full well that we will not completely reach our goal, but also knowing that we will gain much in

the process. On the other hand, if we do not even try to automate intelligence, then we will get nowhere and learn nothing, never realizing the tremendous potential before us. We *expect* that machines will never completely duplicate human intelligence. If we are to try at all, then we must be cognizant of our ultimate goal, and carefully consider the "direction and path of our work." The very essence of intelligence and consciousness begs that we investigate their underlying mechanisms. A natural "next step," so to speak, would be to "automate" those mechanisms. This endeavor should not be taken lightly, however. We provide some words of caution with respect to automating intelligence at the end of this chapter.

The basic function of a computer program is to map a set of input signals to output signals in a meaningful way (note that each neuron performs a similar function at a much lower level). Traditionally, the *meaning* in a computer program is carefully identified during the programming process. The program itself is usually quite rigid, with little or no adaptation or learning built in. However, unless the program incorporates learning, then it will never truly exhibit intelligence. Clever programs may give the appearance of being intelligent, but more often than not programmers would be hard-pressed to defend such a claim. Attempts to automate intelligent information processing have existed for a long time, since before the development of the electronic computer. However, while much progress has been made over the years, computers are still very far from being intelligent. Our belief is that machines will someday contain a high level of intelligence. While we do not know when this may occur, or how the intelligence will be implemented, we do envision this kind of intelligence (i.e., computer intelligence) as being very different from its biological counterpart.

6.1 Describe the "Process"

As we attempt to describe intelligent information processing, we will encounter many problems. These problems stem from the fact that we do not understand the "intelligence process" well enough. *Our ability to model the process depends on our ability to describe the process.* We have tried to describe various aspects of the process throughout this report. We have treated intelligent information processing as a process which transforms data into more useful and significant forms. That is, the process involves transforming raw data into information and then into knowledge as shown by the data spectrum (see Figure 2.1). This process has a direction associated with it, and that direction leads to intelligence. As part of this process, goals must be established. Mechanisms must exist which enable those goals to be achieved. Various approaches must be evaluated, one or more chosen, and energy expended to reach the goals. As part of our approach to modeling intelligent information processing, we try to harness the potential of the process. One manner in which this potential is revealed is in the natural order and organization present in data. We must exploit this order. We must develop and incorporate the transformations necessary to realize and maintain this order in our model. Whether working on specific details or general theory, top-down or bottom-up, a better foundation must first be developed upon which models and applications can be built.

In treating intelligent information processing as a process, we must be able to describe that process. Ideally this would include all aspects of the process. We cannot fully describe the process, but we can try to summarize it. In an earlier report Yaworsky and Vaccaro (1993) stated that:

... data analysis and information processing can be performed in machines using models which perform functions analogous to those in the human brain. The purpose, arguably, is to enable intelligent decisions.

Intelligence requires understanding, which requires communication, which requires language. Understanding involves the ability to learn and to generalize. All of these processes must be realized and accomplished in some kind of network which exploits the inherent components of data. Frequency components have been emphasized in this work. Signals containing high frequency components enter the brain in a parallel fashion. Data gets filtered and focused into various ranges or levels (octaves), resulting in lower and lower frequencies. Harmonic signals get formed as a result of signal interferences in the network. Signals with the strongest harmonics are those which best represent the nature of the data and conform most to the structure of the network. Resulting signals can be used by the network to learn, and can be stored as knowledge in the form of connection weights (memory). The weights represent various forms and kinds of data associations... (which) can be output using some kind of language, resulting in communication, which keeps the process going.

This is one simple example of how the process may be described. Many other descriptions are possible. One thing shown by this description is that we are dealing with an iterative process which, in essence, is concerned with the acquisition of knowledge. As more is learned about the process, more will be asked of it, and more learned. The mere act of focusing on different parts of the process, or focusing on different levels of similar parts, may reveal more information and knowledge. In order to describe the process, we realize that it would be helpful to identify some of the sub-processes involved.

6.2 Identify Sub-Processes

What are the sub-processes of intelligence? With limited knowledge of the entire process, we can try to break it down to further our understanding. From an engineering standpoint, sub-processes lend themselves to implementation more easily than dealing with an entire process all at once. From a human behavior standpoint, we can better understand our thoughts and our actions if we break them down into simpler, more basic components. As we go about our daily business, we perform various intelligent functions, yet we may not know how or why they work. We cannot always explain our behavior. This behavior results from a complex mix of forces and activities in our body and within our brain. These forces affect *mental signals*. The job of the brain, as keeper of intelligence, is to maintain order, harmony, balance and control of these signals. As we stated earlier, the world exhibits natural order, and our understanding of the world depends upon how well we perceive that order. By implementing intelligence as sub-processes which exploit order, we may begin to realize and thus model the forces of intelligence.

Some of the sub-processes (or functions) of intelligence discussed in this report have included goals, decisions, learning, generalization, transformations, filtering, focusing and common sense. Others not well-addressed include perception, attention, awareness, emotions, feelings, thoughts, ideas, understanding, control, communication and language. This is not an all-inclusive list, and the items in the list are not in any special order. No attempt has been made to group them according to importance or combine them according to scale (octave). We mention this only to give an indication of how complex the task of automating intelligence is. Unfortunately, we must defer analysis of many of these topics to another time. We provide a brief description of some of the sub-processes of intelligence in the form of *guidelines* later in this chapter. We

also point out that, in the end, many disciplines will necessarily contribute to the modeling of these sub-processes. A partial list of these disciplines includes artificial intelligence, cognitive science, psychology, biology, neurology, chemistry, physics, mathematics and dynamical systems theory.

6.3 Model Key Components (the "Main Points")

In dealing with a subject as complex as intelligent information processing, one guideline we use to try and keep things in perspective is to "keep it simple, make it useful." This implies that we ought to concentrate on modeling key components of the process, but not necessarily at the expense of overall functionality. One of the key components of intelligent information processing involves the processing of *main points* of thought. Our belief is that the tremendous importance of *main points* with respect to the processes of intelligence must be recognized and exploited. Next we will discuss *main points* in relation to resonance, associations, concepts, context and meaning.

Associations are input-output mappings which represent "cause and effect" relationships. Associations relate input signals to output signals, or stimulus to response. We use stimulus and response here in terms of mental/signal activity, and not in the motor/physical sense. Many associations may be grouped together to form concepts. As concepts form, they naturally inherit relational properties from their constituent associations. Together these properties define the concept, but they also contain a special kind of association, one which involves *context*. That is, when associations and concepts form, references and ranges form along with them, relating signal associations and concepts to appropriate contexts. Context is a function of scale, providing a frame of reference, so to speak, which links a concept or association to an object, event or situation. It is through the activation of concepts (associations) in conjunction with their relevant contexts that signals take on meaning.

Since a concept is composed of many associations, a concept may be considered a set of signals which forms a distribution. Each distribution of signals would contain characteristic features such as reference and range. For a given distribution, the reference may represent an aggregate main point, and the range may reflect the scale or context. The characteristic nature of each distribution would depend on its constituent signals. These signals have properties which are a function of the interaction of basic signal components. The overall process is dynamic, and the nature of the resulting signal activity involves much interaction and interference. For any complex network, the possible number of signal combinations is tremendous. In the brain, this number is seemingly infinite (literally mind-boggling).

Intelligent communication involves getting main points across from sender to receiver. The communication process involves the transmission of signals which have a natural hierarchy built into them. The existence of a hierarchy within the signal structure of the communication process gives an indication that a similar kind of hierarchy is built into the thought process as well. This hierarchy involves such things as associations, concepts and messages which are built up or transformed from basic signal components (i.e., frequency, phase and amplitude components). By combining signal components in different ways, various cognitive constructs may result. For example, the formation of perspectives, the existence of abstraction levels, and even understanding itself are all artifacts of signal interaction. Thoughts are considered to be an orderly arrangement of signal associations consisting of main points, supporting points and relevant contexts. Since thoughts and communication both involve the expression of main points, we

acknowledge their existence and look for ways to model them. This may lead to methods for modeling the processes of thought.

We mentioned in an earlier chapter that signals may resonate in the brain. This kind of resonance is considered beneficial, and not necessarily destructive as may be the case in other situations involving resonance. Here we are talking about resonant activity involving mental signals. This refers to a special kind of resonance in which the response to forces acting on signals in the brain are characterized by intense or enriched vibrations. Resonance results from the natural reinforcement of fundamental modes of vibration. In this case the vibrations involve signals which comprise *intelligence*. These signals interact with each other at many levels, from simple and basic to more complex and abstract. At higher levels of abstraction, many signals may combine to form what we call *resonant concepts*. Each resonant concept has a main point, along with many constituent points, whose associations contribute to the concept as a whole. The resulting signal activity exhibits a natural order or harmony. Depending on the nature of signal activity present (i.e., the circumstances), various aspects of concepts and contexts may be emphasized or lessened. Resonant concepts, in conjunction with appropriate contexts, allow signals to take on meaning, and are thus considered significant components of thought.

Another important aspect of resonant concepts involves direction. That is, we must consider where the resonant concepts are going. Perhaps we allude to this when we use the phrase "train of thought." Where the process is going and how it might get there are significant aspects of the overall *process*. However, while resonant signal activity may be a characteristic feature of signal interaction, resonance is not necessarily a *given* feature. The main point of a concept may or may not be exhibited by resonance. If main points are not well-formed, or if they are buried in a morass of signal activity, signal resonance may not occur. In certain situations, we may not come to fully "realize" the underlying main points. One example of this is confusion. This occurs due to signal interference during which incoming signals are not in accordance with stored signals or any "expectations" resulting from these stored signals. The result is a lack of signal harmony. Resonance, on the other hand, may indicate that a basic understanding, or a so-called harmony of thought, is present. Understanding involves the recognition and realization of underlying order and organization of signals present at a particular time.

In many situations we may never fully realize (understand) what others think or do. We may not even understand certain things that we do. Perhaps in order to understand, we must establish new connections in the brain, thereby enabling appropriate associations and responses to occur. This is the essence of learning. Concepts and contexts are stored (formed) during the learning process, and may be recalled or activated by appropriate signal activity. However, as alluded to earlier, so many combinations of signals are possible that it may take much effort to organize and/or activate a particular sequence of signals. These sequences, or patterns, tend to be quite dynamic in nature. A slightly different set of input patterns can easily result in different signal activity, possibly changing a main point or context, which can drastically affect meaning. This illustrates the complexity of the process and also the sensitive nature of the actual "state of mind" with respect to learning, communication, interpretation and understanding.

Intelligence consists of many complex concepts supported by many complex processes. Forces exist inside the brain (e.g., electrochemical, electromagnetic) which cause various kinds of mental phenomena to occur. For instance, the creation of goals, the making of decisions and the storing of knowledge are ways in which these forces can manifest themselves. This activity leads to various kinds of behavior. One characteristic feature of this behavior is the iterative process called knowledge acquisition. That is,

knowledge is used (in more or less a self-fulfilling process) to gain more knowledge. We are all blessed with intelligence, yet we do not exactly know how it works. We aren't able to see what the "big picture" of intelligence actually looks like. As we attempt to automate intelligence, we must inevitably try to develop this "big picture."

6.4 What Really Matters

A *main point* is a key component of intelligent signal activity. The physical nature of many signals together can be thought of as forming a distribution. A common shape for a distribution is the bell-shaped or normal curve. The peak of the normal curve could be thought of as representing the main point of a distribution. We have suggested that a distribution could be used to physically represent a concept and its constituent associations. But intelligence involves more than the representation and manipulation of main points. At a higher level, or perhaps at a different level of abstraction, we must consider *what really matters*. For the most part this means considering the relative significance of a main point with respect to an appropriate context. This also involves the interaction of goals and circumstances.

To accomplish a goal, main points are typically considered in conjunction with what is going on in the nearby environment. In one sense, *what really matters* depends on the combined effects of the actual "state of mind" (internal signal activity) and the current environment (external signal activity). *What really matters now* is a slightly different phrase which emphasizes a form of short term goal. Sometimes *what really matters now* may not matter much in the near future. This is because as time passes, circumstances may change, or a person's "state of mind" may change. The time-sensitive nature of short-term goals in particular is indicative of the dynamic and time-sensitive nature of information in general. Oftentimes if we look past the short-term nature of a situation we can see a more stable form of *what really matters*. What we are implying here is that goals are a driving force which have a significant impact on us in terms of what really matters over time. Just as working towards achieving a goal is a dynamic process, the process of determining *what really matters* is also dynamic. As we work towards a goal, we must make relevant decisions by considering main points within appropriate contexts. Together this process constitutes *what really matters*. To be able to model this process, we must be able to identify and represent main points and also be able to process the main points within proper contexts. Goals must be modeled as part of the *process*, which necessarily must involve mechanisms for achieving those goals.

When considering *what really matters*, terms, meanings, concepts, goals, context and perspective are important. This indicates the relative nature of *what really matters*. That is, *what really matters* is dependent upon goals and expectations, on knowledge and perspective, and on the meaning of concepts within appropriate contexts. All this constitutes what can be termed the *relative truth* held by a person at a particular moment in time. This may be summarized by the phrase "what really matters now." The slightly different phrase "what matters now" may actually refer to the dynamic processes which constitute consciousness. As another comment on phraseology here, when we use the phrase "in fact" or "as a matter of fact" we may or may not actually be stating a *fact*. One thing to surmise from all this is that the *absolute truth* is elusive.

The overall process involving *what really matters* is iterative in nature. As we go about considering and determining *what really matters* in a given situation, we cannot help but learn. As we learn, we gain knowledge, which better enables us to determine *what really matters* in the future. This iterative process is not unlike many of the other iterative processes comprising intelligence. One difference in each process most likely

stems from the different *kinds* of inputs used by the different processes. For any given process or sub-process, we can envision certain kinds of parallel inputs being filtered to produce fewer significant outputs. This reduction process involves evaluating tradeoffs and making decisions. The criteria used in this reduction process may be any of a number of things. While the criteria may be different, the actual *underlying process* involved may not be so different. *What really matters* appears to be a fundamental sub-process which feeds understanding, intelligence and consciousness. We do not yet know how to represent this sub-process in a model, but we aim to uncover some of its underlying mechanisms.

As part of the process of deciding *what really matters*, we must be able to identify main points. The aspect of time is important, especially when considered in the context of *what really matters now*. It is interesting to note that a longer term perspective can often shed significant light on a situation even when one is mostly concerned with *what really matters now*. As such, desired end-states (goals) may have a strong influence on things even when they appear to be overshadowed by “more pressing matters.” The “bottom line” is also interesting to consider, since it may indicate “main points of thought” and even contribute (consciously or unconsciously) to future generalizations. One product of the overall process of intelligence is knowledge. Consequently, knowledge feeds other crucial “products” such as intelligence and consciousness. As a result of learning, our knowledge base changes, and these changes have a direct and significant impact on our inner self.

With respect to modeling, we can examine *what really matters* and try to incorporate the mechanisms involved into a working model. We may fail to actually state here *what really matters*, but at least we know that it in and of itself is an important process. As such, one of its related sub-processes involves filtering. We are exposed to so much information each day, but we typically focus on only a small amount of this information at any given time. Depending on the situation, we can direct our attention to a relatively small amount of incoming signal activity, understand it, and learn or make decisions from it. In doing this we must naturally filter out a significant amount of information. As such the process must involve trading off between the *quality* and *quantity* of information. *Meaning* is also important (that is, the relative significance of a main point with respect to an appropriate context). As part of intelligent information processing we must be able to trade off between the *quality* and *quantity* of information at any given time in order to determine *what really matters*.

6.5 Guidelines

As we go about trying to model intelligence, or automate intelligent information processing, we look for proper guidelines to help us along the way. As an ongoing exercise, we have compiled a list of terms which we have used over and over again in the course of our work. Unfortunately (and ironically), due to the nature of the list, it is difficult to organize the terms and give their exact meaning. As such, we feel that it is inappropriate to present the list outright. Instead we offer these guidelines which contain some main themes common among the terms. We hope that in doing this we may shed some light on the underlying principles of intelligence. We know that over time our ideas will have to be improved upon and our list of terms will have to be modified. We offer these guidelines as an initial attempt at describing some of the functionality involved in intelligent information processing.

One important guideline when trying to unravel the mysteries of intelligence is to *keep things in the proper perspective*. A perspective is a point of view, with *the point*

being a kind of reference, and *the view* depending upon scale. Goals are needed to keep things in the proper perspective. The ability to assess the relative significance of matters, or *what really matters*, is also crucial to the process. A *perspective* implies that a direction or angle may be associated with a particular *view*. Also associated with a perspective are limits and the issues of *scale*. Scale is a critical issue when it comes to modeling intelligence. We have proposed the octave rule as a way to help filter and focus data into natural ranges or scales. In conjunction with this, we have provided a unique perspective in the form of the data spectrum. We have used the data spectrum to describe the natural order present in data signals as well as to portray the relationships of data signals as they transition from one form to another.

We have stated more than once that with respect to goals, one must be cognizant of the direction in which the goal-process is going. In a practical sense, it may be beneficial to consider the path and direction of the process, including where it has come from and where it may be going. This is similar to considering the sources (and quality) of information used in communication. In addition, very much related to goals are decisions. The mechanisms used to trade off and weigh the various principles and priorities associated with a particular goal are essential to intelligence. As such, the decision-making process is a fundamental sub-process of intelligence.

Another important guideline regarding the intelligence process involves change. We must be able to model change in a system. Dynamical systems concepts are being developed and applied more and more as a means of representing complex, changing behavior. One thing to remember in the modeling of intelligence is that we must ultimately deal with *signals*, and dynamical systems concepts allow us to describe the behavior of those signals as they interact and transition from one form to another on their way towards intelligence. We encourage attempts to represent and exploit the natural order present in all forms of data. As we try to reproduce the *harmony of the process*, we will have to develop mechanisms for controlling the overall behavior of the model, such as using feedback to maintain stability and overall balance.

Another critical issue in automating intelligent information processing is time. It is extremely difficult to understand how the mechanisms of time play in our mind. We can imagine all kinds of scenarios in which time affects our mental processes, but we are a long way from knowing how to actually model the timing mechanisms of the brain. Important concepts related to time are scale and form. Scale is a kind of boundary which defines the forces acting upon signals present in whatever form at any given time. Time, scale and form play together in such a way as to identify the main points of the process when deciding *what really matters*. When it comes to modeling time, scale and form, however, we most likely will struggle with each concept separately until we are able to integrate them in a more appropriate fashion.

Our next guideline concerns *understanding*. This involves being able to establish the meaning of concepts within appropriate contexts. For any given concept, we must be able to determine a relevant main point as well as an appropriate context. Realizing a concept without a proper context (and vice versa) does little good toward contributing to intelligence. As we strive to consider *what really matters* at any given time, we must be able to identify main points among signal activity, and then associate a range or scale with each main point. A main point may be considered a reference, and a range or scale may serve as a context. Taken together, main points and concepts can be used to make sense out of a given situation. Probabilistic and statistical methods will be useful in representing these kinds of processes.

Another guideline stems from our idea that resonance may prove to be an essential part of intelligent information processing. The interaction of mental signals, with the ensuing constructive and destructive interference which must take place in the brain, may result in a maximum response, a so-called "main point," or the production of dominant modes of signal vibration. These dominant modes of vibration may constitute a resonant response necessary for intelligence to occur. This maximum response, or resonance, may consist of the combination of all signals activated at one time and which are related to a particular concept. The trick is in activating the proper combination of signal associations! Since so many biological sub-processes appear to be iterative in nature, it is not unlikely that critical components (i.e., seeds) of intelligence may exist which, when developed properly, can result in signal resonance. Perhaps these seeds are implanted as associations are formed inside the brain. Exploring the phenomena of resonance may provide insight into the *harmony of signals* necessary for understanding the functions of the brain. The frequency, phase and amplitude components of signals may be useful in representing the physical transformations necessary for intelligent behavior to occur.

The architecture of the brain provides insight into how we might want to construct a model of the brain. Scores of processing elements, parallel connectivity and layered functionality offer useful guidelines for the architecture of our model. Ideally the approach taken should be modular, allowing for the incorporation of future developments. The architecture of a model for intelligence must also support the desired functionality from a practical point of view. This means that, aside from all the different theories and implementations which exist, what works in the end should be taken seriously. To this end, theoreticians and practitioners may learn from each other. Both kinds of work are needed. Models will evolve. In any event, it may help modelers to *envision the process from a signal point of view*.

Another architectural consideration involves the concept of weights, or the storehouse of knowledge, as an essential part of learning. As concepts form, they naturally acquire properties or values by way of some scheme which performs the equivalent of *setting the weights*. Through learning, neurons and their interconnections develop in such a way as to build up a propensity to activate, or make associations, based on whatever signal activity is present. Through neural growth and development, the circuitry is defined, and through learning, the pump is primed, so to speak. Weighting mechanisms form which capture the essence of learning. The synapses in the brain are believed to be responsible for much of this functionality. However, due to the fact that there are so many neurons in the brain, and that each one has so many connections, we do not even consider building a model exactly like the brain. Some say that this is impossible to do, while others point out that we can already model some of the brain's functionality today. Since we are really after the *essence* of intelligence and learning, whatever that may be, we would do well to model some of the functionality of the brain, and not necessarily all of it.

As for the issue of representation, our model must ultimately be able to process many *forms* of data. Whether all the data gets transformed into one form for processing, or if it is processed in many different forms and combined later, any realistic model will have to be able to handle various forms of data. This includes dealing with variability and not just black and white concepts. One way to represent many kinds of data in a unified form is to use pulse coding mechanisms. Proper transformation techniques may be used to get various kinds of data into the form of pulses. The pulses can then be processed in an integrated fashion, allowing for the embedded "information content" of various signal patterns to be maintained while at the same time combining different kinds of signal patterns in complex ways. Oscillatory computation techniques may also be used

for more advanced forms of representation and processing. Binary representation and Boolean logic are very useful for many things, but unfortunately for modeling purposes the real world is analog and largely uncertain. One way to deal with these issues is to use more approximate or probabilistic methods. An important factor in whatever approach taken is to *exploit the similarities and differences of signal components* in whatever form the signals may be in. Modeling mechanisms must certainly take advantage of this most basic feature of intelligent signal activity.

In the design of any system, tradeoffs must be made which involve dealing with all the forces present and resolving those forces in a controlled fashion. One way to approach this is to break the forces down into smaller and smaller entities until they can be dealt with as elementary components. In this fashion, mechanisms which handle fundamental positive and negative forces, perhaps in the form of signal *similarities* and *differences*, can be used in a variety of circumstances. This allows trade-offs to be made among the various pros and cons of a situation, and is analogous to evaluating costs and benefits for many different circumstances. The idea is to use simple processes over and over again which will allow the modeling of some very complex tasks.

With respect to intelligence and understanding, we have mentioned the importance and implications of *meaning* to the processes of intelligence. The realization of meaning occurs when a reference or main point becomes associated with an appropriate framework or context. From these associations, more abstract relationships can be built, which can get quite complex when temporal aspects are taken into consideration. Temporal patterns (sequences of data) are constantly being learned and used by the brain. Consciousness involves the realization of these patterns by interpreting the meaning of them in "real-time." The sooner we realize the relative significance of signal activity present, the more cognizant we appear. In order to automate this process, we must be able to implement the essence of meaning in a computer model. As such the model must be able to recognize the *main points* embedded in communication and evaluate their *significance* within appropriate scales or contexts.

One way to approach implementing meaning is to group information according to content or type, and then employ mechanisms to filter that information into finer and finer categories. At the same time, mechanisms which combine resulting forms of data may be used to produce themes or concepts, which would then get acted upon appropriately within the hierarchical model. Obviously this is a crude description of a complex process, but the bottom line is that filtering and focusing mechanisms are essential to information processing, and signal activities equivalent to attention and awareness are needed to transform many complex signals into fewer, more important ones. One possible scenario is that the brain may use some of its machinery for awareness, discerning the big picture and gleaning general information from lots of data, while it uses a different set of machinery for attentional purposes, allowing it to hone in on specific details. This scenario highlights the effect of many signals entering a network (e.g., in parallel) and producing many fewer but more significant signals. As stated earlier, this goes along with our concept of the data spectrum, and is also analogous to the process a judge must perform when making a decision.

Finally, we remind you again that we are out to model the function of the brain, and not necessarily its structure. While we may borrow from the brain's structure, we cannot borrow the entire thing! And while we may emulate certain functions, we cannot exactly duplicate them. We must remember that for all of our efforts, all the ensuing signal activity leads to something. The overall product is intelligence. Whether we model its mechanisms directly or indirectly, we must remember that we are dealing with

a process. We must also be cognizant of the overall direction in which that process is going.

6.6 A Word of Caution

In considering what is involved in the automation of intelligent information processing, we have examined what the field of artificial intelligence (AI) has to offer society, both good and bad. On the good side, AI makes possible many efficient mechanisms for processing information. Computers provide capabilities unmatched by humans (just as machines ought to do). The vast amount of work in AI has helped advance the state of the art in various computer technologies, while at the same time AI aims to incorporate intelligence into automated information processing. This goal of AI is extremely difficult to achieve, as we have seen, but it is arguably man's destiny to pursue this kind of work. However, as with any kind of effort, progress does not come without a price. Writing, industrialization, and advanced forms of travel are examples of some of the great achievements of mankind. Each of them, though, has its own set of problems. For instance, with writing we have such things as libel and propaganda, with industrialization have come waste and pollution, and advanced forms of travel bring with them more catastrophic accidents and more rapid spread of disease. The point is that, with any kind of technology or capability, care must be taken to ensure its proper use. We must be especially careful when it comes to incorporating "intelligence" into machines, as is the case with AI. On the bad side, AI-related technologies can surely be misused. A big part of intelligence involves making decisions, which carry with them associated responsibilities. The bottom line is that humans are ultimately responsible for their decisions, which includes decisions concerning the *control* of their machines. Over time, machines (hardware and software) will be used to perform more and more functions. These machines must be designed carefully and used responsibly. Humans must maintain authority over machines, exercising caution when relinquishing control to them.

7. Impact On Reliability Science and Electromagnetics Technology

Artificial intelligence and neural networks are enabling technologies which will serve to benefit many disciplines. As these so-called *intelligent technologies* mature, more and more of their benefits will be realized, and applications will follow. Also, entirely new disciplines will likely spring from technologies related to artificial intelligence and neural networks. With today's society firmly planted in the computer age due to advances in microelectronics, computer systems and the networking of systems, all trends indicate that computers will have a tremendous impact on society in the future. This impact is felt strongly even today. But with respect to adding *intelligence* to automated methods, the use and benefits of resulting computer technologies can have an even greater impact.

The giant computer industry, however, despite its enormous progress over the past few decades, has a long way to go when it comes to incorporating *intelligence* into computer systems. For the most part this is not due to limitations of computer hardware or software. The inability to automate intelligence stems from the fact that we do not yet understand enough about the intelligence process to be able to model it. The foundations of intelligence have sufficiently eluded mankind's attempts to discover them ever since the beginning of time. Countless efforts to discover and describe the processes of intelligence have come up short, for whatever reason. Whether examining its inherent mechanisms directly or indirectly, we still have not come up with the proper mathematical equations, or even an adequate description in words, to enable us to capture the essence of intelligence in the form of a model. What hope is there? What kind of impact can this scenario have on reliability and electromagnetics technologies?

As it turns out, there is much hope. The natural processes of intelligence may be elusive, but they exist nonetheless. And no doubt they are *orderly* processes. If we can describe (at least some of) these processes and break them down into essential sub-processes, we will be able to model their functionality using computers. Most if not all the math and physics concepts needed to do this already exist. And today's computer technology is powerful enough to at least support the idea that intelligence can be automated. So what is missing? Better models, for one thing, and the knowledge and understanding needed to build these models! The benefits of computers are well known in our society, and it is only a matter of time before these benefits include various (advanced) forms of intelligence. This chapter addresses how efforts to automate intelligence may benefit from, and also be beneficial to, reliability sciences and electromagnetics technologies.

Before we talk about specific disciplines and their possible applications, we must digress a moment. The extent which a new technology may be applied successfully to a particular domain or discipline depends to a large part on the foundation or scientific footing of that new technology. Artificial intelligence and neural networks (collectively considered "AI" here) have been around for forty years, yet they are still considered new technologies. This is because their acceptance and use has not yet become widespread. Since *fundamental principles* in AI have not yet firmly been established, people may still consider them (relatively) new technologies. Researchers have tried feverishly to put these so-called *intelligent technologies* on firm footing, with limited success. Of course many applications of these technologies exist today, but all in all, the so-called intelligent technologies are not an overwhelming success. Critics argue that AI has not been successful at all, and skeptics say that the goal of automating intelligence is impossible to

achieve. In any event, AI is a very real technology and an important part of the natural evolution of computers.

The core functionality of intelligence has not yet been discovered, and certainly the essence of this functionality has not yet been automated. But try we will. Computers are here to stay, and as we allow them to do more and more for us, it is inevitable that they approach doing so more *intelligently*. One of the many troubling aspects resulting from this is the fact that, from a development point of view, attempts to automate intelligence are often so entwined with real intelligence that it may be difficult for the casual observer to tell the difference between the intelligence of a computer program and that of the programmer, or between the intelligence imbedded in hardware and that of its designer. This distinction is not always clear-cut, and the question of who (or what) performs so-called intelligent functions in AI applications is often difficult to resolve. This situation will likely worsen in the future, as we allow computers to do more and more for us and as AI technologies continue to mature. Computer researchers know all too well how they must painstakingly compensate for what the machine (hardware and software) cannot do with respect to intelligence. Now and in the future, complex issues in AI-related applications may not be well-defined, the consequences of running programs may not be well-known, and the responsibilities for actions taken may be left unaccounted for. As we go about developing our so-called *AI applications*, we must also (collectively) develop certain ground rules for these applications.

7.1 Signal Activity

Many disciplines contribute to the task of automating intelligence, with perhaps as many different approaches as there are research groups. Conventional AI researchers have concentrated heavily on symbolic and other high level aspects of intelligence, while neural network researchers have emphasized approaches based on neurons and their weighted connections. Other disciplines have examined the nature of intelligence from a biological or psychological point of view, while still others have looked at various signal processing techniques. But for the most part, *the basic signals underlying intelligence have not been modeled well enough*. In this context *signals* imply mental signals, but these signals can originate from many different physical sources (and be used by computers also!). While the potential causes of signal activity may be pondered for a long time, and the many possible effects of signal activity within the brain may never exactly be known, the actual physical signals can nonetheless be considered one of the raw materials enabling intelligent behavior. Of course these signals become embodied in a brain and acted upon via a human body, but mental signals are an important piece of the puzzle. We realize that the nature of these signals is abstract, and the interaction of these mental signals with the physical body (i.e., the "mind") are issues of great debate. But we also acknowledge that *signal activity* plays a crucial role in the establishment of intelligence. Consisting of frequency, phase and amplitude, these *signals* represent a fundamental component to the physical aspects of intelligence.

While the significance of signals to processing, computation and the engineering of physical systems has been recognized and used in various disciplines, researchers are far from exploiting *intelligent* signals, and even further from *automating* intelligence as a whole. Today's computers thrive on digital signals, precise timing and explicit instructions. But in reality, the Boolean principles upon which computers are based are an extreme idealization of the natural world. That is, the 1's and 0's inside a digital computer are an extreme form of representing the mostly analog world. In many cases discrete representation does well, but in many other cases it falls short. We cannot exploit the use of (i.e., model) intelligent signals until we better understand them.

Human communication typically involves the interchange of many signals. These signals exist in various forms and carry complex meanings at different levels of abstraction and sophistication. Each instance of communication may involve inherent subtleties or assumptions between sender and receiver. So much is built into normal communication that we would probably be amazed at what must occur (from an engineering standpoint) in order for an "understanding" to be reached when people communicate. Computers do not share the complexities or capabilities of the human brain, nor do they "understand" anything the way we do. Computers must explicitly be told what to do by way of a program. Little by little, these programs may be designed to incorporate some form of "learning." Humans, on the other hand, learn naturally, and their "programming" implicitly involves learning.

The tremendous gap which exists between the highest levels of intelligence and its fundamental signal components is not just a problem for the computer industry. However, the problem seems to be exacerbated in computers by the use of discrete representation, formal logic and certain forms of reasoning. The links between intelligence and signal activity will remain obscure until we can better understand the fundamentals of intelligence and how to exploit basic signal activity in automated models. We all have some understanding (at a high level) of what constitutes intelligence by virtue of using and observing it each day, but we would be hard pressed to describe it. No one has been able to explain the finer details of intelligence and how these details relate to the more general functions of the brain.

Our approach to automating intelligence emphasizes the use of signals whose basic components can be used as building blocks in a hierarchical model. This involves the development of various kinds of signal manipulation techniques stemming from signal activity, including transformations, interference techniques (both positive and negative) and dynamic interactions. This also requires a perspective which shows how lower level signal components may combine to construct higher level functions, leading to the formation of associations, concepts and sequences of events. Our signal-based approach acknowledges that order and organization are a natural part of intelligent signals, and we aim to uncover and exploit that order.

The approach of using basic signal components for various kinds of processing and analysis is not new. Many disciplines have done just that for years. Each discipline has its own way of gathering, representing, processing and interpreting data. To a large extent, the sophistication of any system or discipline depends on how well it handles its data. Intelligent analyses and methods which interpret signal activity help define a discipline, and ultimately these functions contribute to that discipline's foundation. These intelligent functions may be passed on from generation to generation, or they may be learned first-hand. Interestingly enough, the actual signal activity which initiates and contributes to learning can be quite simple, and the signals may pass by quickly, but their effects (i.e., their meaning) can last a lifetime.

A signal may be described as a physical embodiment of a message. The physical nature of a message, with its many signal components, defines the kind of information the message may convey. Communication is considered a necessary function of intelligence. During communication, information is conveyed from source to destination. Being the essence of communication, signals may be perceived as the lifeblood of intelligence. One of the mysteries of this is understanding how relatively simple, basic signals can actually contribute to the formation of something as complex as intelligence.

The signal-based approach to automating intelligence aims to examine the physical nature of the intelligence-phenomenon across many levels. Modeling attempts must begin with a description of the phenomenon. Depending on what level of functionality is being modeled, a mathematical description may suffice, or a more qualitative or textual description may be required. Herein lies one of the major difficulties in attempting to automate intelligence: a detailed description (involving appropriate representation and meaning) is required for conventional automated methods, yet we do not know enough to provide such a detailed description. On the other hand, less precise descriptions may enable the modeling of some intelligent functions, but these are inherently more difficult to automate.

In our efforts to automate intelligence, we acknowledge that many approaches must be investigated, combined and improved upon. No single approach has the market on *the nature of intelligence* at any level. Our emphasis has been on a signal-based approach, which naturally draws from many disciplines. This basic approach simply hasn't been applied well enough to the automation of intelligence. There are many issues to resolve, ranging from general to specific. As we investigate ways to automate intelligence, we can learn from existing techniques in more mature disciplines. Next we will discuss how efforts to automate intelligence may benefit from, and also be beneficial to, reliability sciences and electromagnetics technology.

7.2 Physical and Mathematical Concerns

The physical and mathematical aspects of signal components and their resulting behavior provide a significant contribution to the automation of intelligence. Mathematical methods will necessarily be used to describe the physical nature of signals, and the overall behavior of signal activity must conform to the laws of physics. In time our knowledge and understanding of the intelligence process will progress. The intelligence process is continuous, as knowledge transitions from one generation to the next. Understanding is a more integral part of intelligence than knowledge and represents the implementation of knowledge in an individual for a particular purpose. We should not expect to discover or understand intelligence fully, but to work toward improving what is already known. As we move in the direction of increased understanding, we must be concerned with basic principles in math and physics, and investigate ways in which these concepts may relate to the automation of intelligent information processing.

The reliability sciences and electromagnetics disciplines each have technology concerns related to the automation of intelligence. The study of electromagnetics involves a variety of phenomena and signal activity which act both in and around us. The physical and mathematical aspects of electromagnetic energy have been studied for well over a century, with many significant developments being made. However, as far as the brain is concerned, we know too little about how electromagnetic energy affects internal functions, or how the brain generates and uses its own electromagnetic energy. Yet as Roger Penrose put it, "there is no doubt that electromagnetic phenomena have relevance to the workings of our brains" (Penrose, 1989). In time we will come to know better the significance of electromagnetic effects within the brain, but in the meantime, we go by the premise that the behavior of the brain and that of other electromagnetic systems have more in common than we now realize. Advancements in one discipline can lead to developments in another.

Basic principles in the reliability sciences involve different but no less important phenomena. Reliability involves a study of the natural forces acting upon a system, and represents a measure of how that system will operate under certain conditions over a

specified period of time. The science of reliability is rooted in probability and statistics. Mathematical modeling and analyses can indicate which reliability parameters most directly impact system performance, leading to better design tradeoffs, more cost effective system development and improved quality. All physical systems involve actions and interactions which have consequences. At some point these consequences may be perceived as signals, the nature of which provides an indication of system behavior. The better we can model this behavior, the better we can predict system performance.

Fundamental similarities exist between reliability, electromagnetics and the so-called intelligent technologies at various levels. At a high level, each discipline must deal with vast amounts of data (signal activity). This data exists in various forms, and its presence ultimately affects technology-related decisions. The implication here is that the decisions are "intelligent" and that the data has a direct impact on the decisions. In reality, this is often not the case, with complex, indirect and often unknown relationships existing among data and any conclusions drawn from them. The decision-making process consists of many lesser processes. Interactions and combinations of data signals present at any one time contribute to the hierarchy of intelligence. The organization and behavior of these interactions create information, which may then be transformed to knowledge and used for the purposes of intelligence.

At lower levels of analysis, commonality exists by virtue of similar methods for handling signal activity. This includes representing and processing signal associations and patterns using various tools and techniques of the trade. In reliability, cause and effect relationships are represented using principles which are rooted in physics and mathematics, especially in probability and statistics. These include such considerations as averaging, approximations, assumptions, confidence intervals, combinations and uncertainty. An overall goal of reliability is to ensure that system performance meets expectations. Efforts are aimed at achieving quality systems at affordable cost. In electromagnetic technologies, the main concern is to ensure that electromagnetic behavior of a system is acceptable, and that electromagnetic performance meets expectations. This not only involves the investigation of electromagnetic energy present within a system, but also the coupling and interference of signals in the surrounding environment. Resulting signal interactions can be very complex. Yet electromagnetic energy is a highly *organized* form of energy. This organization involves an "information" component as discussed in Chapter 1. Insight may be gained from the study of electromagnetics which will aid in the automation of intelligence. Likewise, efforts to automate intelligence will spur new ideas involving electromagnetic effects and related system behavior.

From a physical point of view, intelligent information processing, reliability and electromagnetics each involve physical signals entering and interacting within a complex system. From a mathematical point of view, equations can be developed which represent the nature and behavior of these signals. This drastic simplification does not mean that automating this kind of signal activity is a straightforward process. The process has proven to be anything but straightforward. However, signal components (frequency, phase and amplitude) contribute to the natural harmony of a system, and beginning with basic principles of signal components we can build models which reflect the natural hierarchy and organization of an entire system. Signal harmony is a by-product, or *effect*, of a system. The root *cause* of the harmony is the interaction of fundamental forces (in this case signal components) naturally present in a system. The most significant point of all this is that signals in all forms contain orderly components (e.g., information). We aim to develop and exploit this fact.

7.3 Manual Tools of the Trade

Rome Laboratory (formerly Rome Air Development Center) supports a wide variety of technology research and development in electromagnetics and the reliability sciences. As part of a continuing commitment to satisfy Air Force mission requirements, electromagnetics and reliability are two of the enabling technologies developed and coordinated at Rome Laboratory on behalf of the Air Force, and in some cases, the entire Department of Defense. Various resources exist in the form of tools, techniques, services and of course the expertise needed to evaluate the electromagnetic performance and reliability requirements of electronic systems. Capabilities include unique design, analysis, modeling, simulation and test facilities at various locations in New York and Massachusetts.

Historically the tools of the trade in reliability and electromagnetics have been developed by hand and applied manually to mission problems. With these disciplines deeply rooted in the mathematical and physical sciences, tasks often involved calculations which were tedious and labor-intensive. Whenever possible, these tasks were initially performed with the help of machines (e.g., calculating devices or electrical test equipment). These disciplines have developed and matured over the years to the point where manual tools and techniques have given way to sophisticated automated methods. The foundation of these disciplines, however, was laid by the development and use of manual methods. Today's electronic technologies, from transistors to computers to the internet, have planted us firmly in the information age. However, we should not forget that the precursors to these technological advancements were manual methods, having helped establish the reliability and electromagnetics disciplines into substantial and essential technology areas.

In the electromagnetics domain, work at Rome Laboratory has concentrated on the generation, transmission, detection and interaction of electromagnetic energy within Air Force systems. Concerns have ranged from electromagnetic effects of microelectronic devices to the performance of entire platforms. This includes internal electromagnetic behavior as well as coupling effects due to signals from nearby antennas. A related concern deals with system behavior in environments which are electromagnetically rich, whatever the source of signals. As one of the fundamental forces in nature, electromagnetic behavior is a crucial concern in the development of electronic equipment.

In the reliability sciences, emphasis is on the design, analysis, fabrication, testing and specification of electronic components and systems (see, for example, Morris and MacDiarmid et al., 1995). Reliability modeling includes allocation, correlation, assessment, classification, optimization, recognition, estimation, diagnosis and prediction. Quality is a big part of the ultimate goal to ensure adequate system performance at affordable cost over a system's lifetime. Continuous improvement of relevant processes is an important part of the quality program. This includes monitoring, identifying and controlling variations of key parameters in order to optimize the underlying processes. There can be no doubt that reliability is a critical discipline -- it must be considered sooner or later, and from a cost point of view, sooner is better.

As part of every engineering application, decisions must be made with respect to system performance and reliability. Cost is always a big factor, but certainly not the only one. Many other considerations must be made and tradeoffs continually evaluated to support system-related decisions. Because of the difficulties involved with understanding and evaluating decision criteria, the need exists for better data collection, representation,

analysis, and modeling techniques. This implies the need for automation. While manual tasks and new ideas will always be necessary, more and more emphasis is being made, and energy expended, on efforts to automate these processes. To help improve reliability modeling for electronic systems, and to help evaluate the behavior of complex electromagnetic systems, computers are being used to facilitate the investigation and application of these technologies.

7.4 Automating Tools of the Trade

Given the widespread availability of computers today, more and more tasks in virtually every discipline are being automated. This is driven by the natural tendency to perform job functions as efficiently as possible. Automation provides the ability to perform tasks quicker, easier and more accurately. Computers also enable many tasks to be performed which were previously impractical or just not possible. In the reliability sciences and electromagnetics disciplines, many tools and techniques have been developed over the years which provide a variety of data analysis and information processing capabilities. Advances in modeling and simulation techniques have drastically changed the way designers, engineers and analysts do their jobs.

Many of the benefits associated with automating tools of the trade in reliability and electromagnetics include the ability to design, test, analyze and validate system performance in a simulated or synthetic environment. This may occur in conjunction with actual development, in parallel with it, or it may involve theoretical research accomplished in lieu of physical system development. Computer simulations produce a model which by definition is not the same as the real thing, but scientists and engineers are finding that computer modeling can be a viable alternative to actual development. Whether it involves theory or practice, investigating new techniques or improving existing ones, computers offer a substantial mechanism for performing tasks in today's workplace.

Rome Laboratory has been involved in the development of many automated reliability tools over the years. Tasks performed by software packages run the gamut in data analysis (Fuqua, 1993). Actually, the functions performed in software are limited only by the constraints which exist due to the state of the art in computer modeling and existing theory. For example, statistical techniques not only include quantitative methods but may also include qualitative ones, such as those available through approximate reasoning techniques (e.g., fuzzy logic). Automated techniques also provide the ability to filter through enormous amounts of data now available from sensors, diagnostic circuitry and networks of databases. The limiting factor in analysis today is no longer the lack of data or processing power, but the inability to model appropriate processes which translate data into something more useful.

Electromagnetics work at Rome Laboratory mainly focuses on the performance of antennas used in Air Force systems. The behavior of electromagnetic signals in the form of antenna patterns are evaluated through various means to ensure that the interference or coupling of electromagnetic energy is such that it conforms to system specifications. Models have been developed which help investigate electromagnetic phenomena and enable the analysis of electromagnetic fields in Air Force systems (Siarkiewicz, 1987).

Ironically, one realization of the work performed at Rome Laboratory has indicated that while certain benefits may be gained from the automation of tools and techniques, other problems may be introduced *because of* the automation. For instance, each automated technique requires formal specification and formatting of software to

ensure its proper use and operation. Also, hardware which can execute the software must be readily available. While this may not be a problem in some applications (such as for standalone packages), conflicts may arise as modifications are required or as automated packages must interact and share data across platforms. The current trend in information processing and concurrent engineering involves the integration of resources which may be physically separate. The issues of networking and combining diverse resources are becoming increasingly important, and the success of future systems development depends on methods which will alleviate problems resulting from this integration and automation.

Computerized models are designed to process data in an efficient manner. The speed and precision of current computer technology drastically outperforms human capabilities for certain kinds of operations. But in many other instances, such as image recognition, dynamic decision making and the formation of ideas, the computer is just no match for human capabilities. Together, however, the possibilities are virtually limitless. Humans provide the intelligence, and computers provide the means to automate. We live in a data driven world. That is, the world is an orderly one, and data (signal activity in one form or another) is the key to revealing this order. The ultimate use for data is to support the processes of intelligence, in one way or another. We work toward implementing these processes in the form of automated tools and techniques.

7.5 Adding Intelligence to Automated Tools

We have mentioned many times in this report that while automated methods are becoming more and more widespread, adding intelligence to these automated methods will not be easy. This goes for the automated tools and techniques of reliability and electromagnetics. Each discipline will be impacted differently by the automation of intelligence, and each will have to deal with advancements in computer technology as it sees fit. In the past, AI has had a relatively small impact in the reliability sciences and electromagnetic technologies. Various efforts at Rome Laboratory have investigated using AI technologies to facilitate or augment the implementation of certain tasks, especially diagnostics and testing (Cooper et al., 1989 and 1991, and Broadwater et al., 1995). The application of AI technologies is also being investigated in attempts to simplify the electromagnetic modeling and simulation process (Drozd et al., 1996). However, due to the subjective nature of modeling, the lack of maturity in AI techniques and the overall complexity of relevant tasks in general, AI applications have achieved limited success. As we learn more about the underlying processes of intelligence, and as modeling techniques improve, we will become better able to realize the potential benefits resulting from the integration of these technologies. Due to fundamental similarities between reliability science, electromagnetic theory and the processes of intelligence, mutual benefits will be gained by advancements in each discipline.

Reliability and electromagnetics may benefit from the automation of intelligence in various ways. The statistical and probabilistic principles which form the basis of reliability sciences may be enhanced by data analysis techniques fashioned after those occurring in the brain. The ability to adapt to a changing environment and to represent and process complex data relationships may inspire new ideas and enable the development of better statistical models. Likewise, electromagnetic modeling will benefit from "smarter" techniques developed to simplify or otherwise ease the computational burdens associated with electromagnetic modeling. Processing capabilities fashioned after signal activity in the brain may provide insight into problems associated with identifying, transforming and combining pertinent signal information. For instance, rules of thumb and heuristic techniques which have been developed over the years may be encoded into rules or developed into other kinds of (intelligent) algorithms.

We may also glean knowledge from fundamentals in electromagnetic technology as they relate to the interaction and transformation of data in various forms as portrayed by the data spectrum (see Chapter 2).

We have discussed various topics which are related to the automation of intelligence in Chapter 5, such as goals, decisions, common sense and the importance of time, scale and form. Future research must address these and many other aspects of intelligence. But in the end, software programs must include the ability to learn in order to be considered intelligent. Learning involves being able to exploit differences and handle variations in data. This functionality may be accomplished in different ways and is in fact difficult enough to achieve to warrant the collaboration of researchers in diverse technical areas.

As we work toward automating intelligent information processing, we must acknowledge the overall goal, but also realize the general direction in which the work is proceeding. In this report we have not provided details on how to automate intelligence, but we have proposed an approach which describes some ideas to pursue and some feasible paths to take. We have not established a sufficient theoretical foundation, yet we have identified simple components of signal activity and have hinted at how these may contribute to higher levels of intelligence. Along the way we must learn to recognize, both in "real life" and in our models, that significant associations exist among the many lesser signal components present at any one time. We need to model the process which can identify such associations.

Benefits stemming from the automation of intelligence will eventually be realized, some being generic in nature while others applying directly to only a few disciplines. Much of the necessary physical and mathematical constructs needed to automate intelligence exists today. What we need are better models. We have searched for a better framework, and have offered a different perspective on the issue of automating intelligence. As is true for any technology, if it is applied too early (i.e., before it is on a solid foundation), then the application will not last. That is not to say that there are no benefits in trying! On the contrary, there are many reasons for applying technologies as early as possible. But in the long run, a good foundation is a must. In any event, progress will occur over time, with many different researchers working to realize this progress.

In this report we have described a particular kind of process. This process involves many technical disciplines and encompasses a natural progression. The progression begins with certain natural forces which produce interesting behavior whose phenomena may be physically observed and mathematically described. This behavior may be characterized in models which are developed to describe and control the behavior. The models can then be used to help solve problems associated with the resulting behavior. Over the years many such models have been developed by hand and manipulated manually in virtually every discipline. Lately many of these models have become automated. Eventually these models may come to incorporate some form of intelligence as part of their nature. An underlying premise of our work is that this progression, this process, must first be *mentally* realized before it can be *physically* described in the form of a model. Ultimately this means that we must realize the essence of intelligence. Once we "perceive" this enormous stumbling block, we can begin to transform it into smaller stepping stones.

8. Summary

We have examined the overall process of intelligent information processing and tried to break it down into basic sub-processes. The goal is to automate some of these processes and apply them to domain-specific problems. We have had to be open to many different ideas, concepts and developments in order to do this. We realize that we fall far short of finishing the task. We also realize that we must be open to change, for the processes we are after are so dependent upon change. As we consider the many possibilities of signals interacting in the confines of our model, we must figure ways to orchestrate and direct their activity. In other words, we need to figure out how the signals play together. In our attempts to achieve this so-called harmony in our model, we know that we have optimized versions of similar processes in our own brains. The time is drawing nearer when we may automate some of these processes.

We take so much for granted when it comes to intelligent information processing. For instance, how the powerful and elusive concept of *information* may lead to intelligence is largely a mystery. Before we can automate something, and do it well, we really have to understand the basic operations and functions involved. We have provided a unique perspective on our approach to automating intelligent information processing. We have come to certain realizations concerning the processes involved, even though our realizations come in the form of a conceptual model. The most significant realization is that a *process* is at work in intelligent information processing. We have discussed how order and change are fundamental to this process. Another important consideration involves what the process may lead to, or where it is going. We have used the concept of the data spectrum to show how data may be transformed into knowledge. We have also used the octave rule to illustrate how filtering, focusing and scaling are part of the process. With respect to modeling and artificial intelligence, we have emphasized the approach of manipulating signals rather than the more conventional AI approach of manipulating symbols. Other modeling concerns involve the use of pulses, oscillators and "state attractors" to model dynamic behavior. Intelligent information processing is most definitely a dynamic process. As for reliability and electromagnetics, we have seen how these disciplines both contribute to and benefit from the automation of intelligent information processing. From the reliability sciences have come many theories and techniques for modeling and analyzing different forms of data, with emphasis on probabilistic and statistical methods. In electromagnetics, models characterize signal activity and interaction with respect to the electromagnetic spectrum. Both of these disciplines are believed to be major, related, and largely untapped sources of information applicable to automating intelligent information processing.

As we strive to develop appropriate models, we try to keep our model simple yet useful. We know that the functionality we are after is quite complex, but we hope that our ideas and perspective offer a unique contribution to the task at hand. Our description has been simple and quite general, but one thing we have learned is that general concepts can be very powerful and useful. If done properly, generalization allows the synthesis of many simple ideas and concepts into fewer, more significant ones. As such, this report is only an introduction. As the process becomes better defined, it will require continuous improvement along with the development of many necessary details. We know we have raised many questions and given too few answers. The motivation for this work is that the benefits will far outweigh the effort. This description is part of our quest for knowledge about a process we use every day of our life, yet know too little about. We hope that the work described here will help enable the development and realization of this evolving and elusive process into workable models. Our contribution is small. We know we have only scratched the surface. But we believe this is a step in the right direction.

9. Final Remarks

Researchers in artificial intelligence have always faced difficult problems, working in largely uncharted territory. One of the more sticky problems has been this: fundamental principles, or underlying concepts, describing and defining intelligence and intelligent information processing do not exist in a tangible form! This undermines efforts to "automate" them. A better foundation and framework are needed. With a solid foundation, robust models can be built. To see the framework, one must have a view of the big picture. This involves having an appropriate perspective. It is along these lines that we have presented our work here.

Our general goal is to enable more efficient (i.e., intelligent) processing of information. We aim to develop better modeling tools and techniques using mathematical transforms, advanced representation and appropriate processing techniques. We must break the "process" down into fundamental principles and then build applications upon these principles. This will occupy the efforts of researchers for many years to come. The important point here is that we work "toward" a solid, fundamental goal, one that, although not easy to achieve, will not change much and is worthy of our time and energy. To that end, an overriding goal is to identify and describe the phenomena involving intelligence, and to sufficiently model that behavior using computer technology.

We recommend, as a minimum, a dual approach toward automating intelligent information processing. One approach involves looking at the entire picture. We need to understand the intelligence process in general before we can sufficiently model it. This not only provides direction and balance to the development of a model, but also helps provide a much needed description of the intelligence process. Even at this high level, though, we must break the process down and examine the kinds of functions needed. For instance, scaling, transformations, analysis, synthesis and interaction of signals are fundamental to the process. As we work to understand and describe these sub-processes, we must develop a functional model which incorporates the necessary general principles.

The second part of the dual approach involves developing fundamental components for the model. This requires that tools and techniques be developed in detail, ones which allow the desired functionality to be implemented. We must be able to build our model using elementary components, detailed equations and specific techniques. This collection of fine details will dictate how the physical model will operate. The fundamental behavior of the model will initially be implemented in the form of computer simulation. Ultimately the two approaches (top-down and bottom-up) must be combined. The resulting model will consist of detailed techniques which are compatible with high level functional blocks, forming a cohesive hierarchical network.

As this dual approach achieves results, they can be used to benefit acquisition cycle design and decision-making processes in many ways. A typical set of tasks in phenomenological modeling and simulation would be to gather and filter data, identify significant parameters, establish limits, recognize trends, match patterns, perform trade-off analyses, make decisions and achieve goals. These tasks are part of the problem-solving process. We perform these tasks naturally. One could even say that we learn to do them "automatically." However, while computers do everything "automatically," they can hardly learn to do any of these tasks naturally. We aim to reveal portions of the intelligence process and incorporate useful techniques into computational models.

10. References

- R. Abraham and C. Shaw, 1989, *Dynamics - The Geometry of Behavior, Part 1: Periodic Behavior*, Aerial Press Inc., Santa Cruz, CA, pp. 40-45.
- M. A. Arbib, Editor, 1995, *The Handbook of Brain Theory and Neural Networks*, MIT Press, Cambridge, MA.
- M. S. Bartlett, 1975, *Probability, Statistics and Time*, Chapman and Hall Ltd., London.
- L. Bookman, 1994, *Trajectories Through Knowledge Space - A Dynamic Framework for Machine Comprehension*, Kluwer Academic Publishers, Boston.
- * S. Broadwater, D. Lunde and B. Havlichek, 1995, *Intelligent Built-In Test and Stress Measurement*, Rome Laboratory Report RL-TR-95-43.
- J. D. Bronzino, Editor, 1995, *The Biomedical Engineering Handbook*, CRC Press, IEEE Press.
- J. Caroli, F. Fieldson, J. Hewitt and M. Plano, 1989, *RADC Fault Tolerant System Reliability Evaluation Facility*, Rome Air Development Center Report RADC-TR-89-165.
- J. Caroli, G. Lyne and J. Vaccaro, April 1990, *RADC Automated R/M/T Design Tools, 6th RAMCAD Technical Interchange Meeting*, Society of Logistics Engineers, Huntsville, AL.
- M. Caudill and C. Butler, 1990, *Naturally Intelligent Systems*, MIT Press, Cambridge, MA, p. 155.
- * C. Cooper, K. Haller, J. Zbztnewski, K. Anderson and J. Matson, 1989, *Smart BIT-2*, Rome Air Development Center Report RADC-TR-89-277.
- C. Cooper, K. A. Haller, V. G. Zourides, P. Skeberdis and W. Gibson, 1991, *Smart BIT/TSMD Integration*, Rome Laboratory Report RL-TR-91-353.
- E. Cox, 1994, *The Fuzzy Systems Handbook*, AP Professional, Boston.
- D. Crevier, 1993, *AI; The Tumultuous History of the Search for Artificial Intelligence*, Basic Books, NY, pp. 240, 258.
- F. Crick and C. Koch, 1992, The Problem of Consciousness, *Scientific American*, Vol. 267, No. 3:152-159.
- D. C. Dennett, 1991, *Consciousness Explained*, Little, Brown and Co., Boston.
- R. L. Devaney, 1993, *Chaos, Fractals & Non-Linear Dynamics*, IEEE Videoconference, 25 March 1993.
- A. L. S. Drozd, T. W. Blocher, V. K. C. Choo and K. R. Siarkiewicz, March 1996, The Intelligent Computational Electromagnetics Expert System (ICEMES), *Applied*

Computational Electromagnetics, The Applied Computational Electromagnetics Society, Monterey, CA, pp. 1158-1165.

L. Fu, 1994, *Neural Networks in Computer Intelligence*, McGraw-Hill, NY.

N. Fuqua, 1993, *Reliability & Maintainability Software Tools*, Reliability Analysis Center, Rome, NY.

J. Giarratano and G. Riley, 1989, *Expert Systems, Principles and Programming*, PWS-Kent Publishing Co., Boston.

J. Gleick, 1987, *Chaos - Making a New Science*, Penguin Books, NY, pp. 221-232.

S. Grossberg, 1995, The Attentive Brain, *American Scientist*, Vol. 83, No. 5:438-449.

S. Haykin and X. B. Li, 1995, Detection of Signals in Chaos, *Proceedings of the IEEE*, Vol. 83, No. 1:95-122.

F. Hlawatsch and G. F. Boudreaux-Bartels, 1992, Linear and Quadratic Time-Frequency Signal Representations, *IEEE SP Magazine*, Vol. 9, No. 2:21-67.

J. Horgan, 1994, Can Science Explain Consciousness?, *Scientific American*, Vol. 271, No. 1:88-94.

J. L. Johnson, 1994, Pulse-Coupled Neural Nets: Translation, Rotation, Scale, Distortion, and Intensity Signal Invariance for Images, *Applied Optics*, Vol. 33, No. 26:6239-6253.

E. R. Kandel and J. H. Schwartz, Editors, 1985, *Principles of Neural Science*, 2nd ed., Elsevier, NY, pp. 49-57, 75-107, 727-770.

M. J. Kiemele and S. R. Schmidt, 1990, *Basic Statistics: Tools for Continuous Improvement*, Air Academy Press, Colorado Springs, CO.

A. N. Kolmogorov, 1956, *Foundations of the Theory of Probability*, Chelsea Publishing Company, NY.

H. Lahore, 1984, *Artificial Intelligence Applications to Testability*, Rome Air Development Center Report RADC-TR-84-203.

F. Machlup and U. Mansfield, Editors, 1983, *The Study of Information*, John Wiley & Sons, Inc., NY.

S. Morris, P. MacDiarmid et al., 1995, *Reliability Toolkit: Commercial Practices Edition*, Reliability Analysis Center, Rome, NY.

J. D. Murray, 1989, *Mathematical Biology*, Springer-Verlag, NY.

H. Pashler, 1993, Doing Two Things at the Same Time, *American Scientist*, Vol. 81, No. 1:48-55.

R. Penrose, 1989, *The Emperor's New Mind - Concerning Computers, Minds, and the Laws of Physics*, Oxford University Press, NY, pp. 124-129, 150, 444.

T. Rossing, 1990, *The Science of Sound*, 2nd ed., Addison-Wesley Publishing Company, Inc., Reading, MA, pp. 33, 127.

C. E. Shannon and W. Weaver, 1949, *The Mathematical Theory of Communication*, University of Illinois Press, Urbana.

K. R. Siarkiewicz, May 1987, GEMACS - An Executive Summary, *Applied Computational Electromagnetics Newsletter*, The Applied Computational Electromagnetics Society, Monterey, CA, pp. 124-136.

C. Skarda and W. Freeman, 1987, How Brains Make Chaos in Order to Make Sense of the World, *Behavioral and Brain Sciences*, 10:161-195.

T. Stonier, 1990, *Information and the Internal Structure of the Universe - An Exploration into Information Physics*, Springer-Verlag, London, p. 111.

S. Strogatz and I. Stewart, 1993, Coupled Oscillators and Biological Synchronization, *Scientific American*, Vol. 269, No. 6:102-109.

N. Wiener, 1967, *The Human Use of Human Beings - Cybernetics and Society*, Avon Books, NY, pp. 50, 277-278.

P. Yaworsky and J. Vaccaro, 1993, *Neural Networks, Reliability and Data Analysis*, Rome Laboratory Report RL-TR-93-5, pp. 8, 39.

M. Zaghloul, J. Meador and R. Newcomb, Editors, 1994, *Silicon Implementations of Pulse Coded Neural Networks*, Kluwer Academic Press, Boston.

* RL-TR-95-43 - Distribution limited to DOD & DOD contractors only;
Critical Technology.

* RADC-TR-89-277 - Distribution limited to US Gov't agencies and their
contractors; Critical Technology.

MISSION
OF
ROME LABORATORY

Mission. The mission of Rome Laboratory is to advance the science and technologies of command, control, communications and intelligence and to transition them into systems to meet customer needs. To achieve this, Rome Lab:

- a. Conducts vigorous research, development and test programs in all applicable technologies;
- b. Transitions technology to current and future systems to improve operational capability, readiness, and supportability;
- c. Provides a full range of technical support to Air Force Materiel Command product centers and other Air Force organizations;
- d. Promotes transfer of technology to the private sector;
- e. Maintains leading edge technological expertise in the areas of surveillance, communications, command and control, intelligence, reliability science, electro-magnetic technology, photonics, signal processing, and computational science.

The thrust areas of technical competence include: Surveillance, Communications, Command and Control, Intelligence, Signal Processing, Computer Science and Technology, Electromagnetic Technology, Photonics and Reliability Sciences.